

The University of Sydney

Copyright in relation to this thesis*

Under the Copyright Act 1968 (several provision of which are referred to below), this thesis must be used only under the normal conditions of scholarly fair dealing for the purposes of research, criticism or review. In particular no results or conclusions should be extracted from it, nor should it be copied or closely paraphrased in whole or in part without the written consent of the author. Proper written acknowledgement should be made for any assistance obtained from this thesis.

Under Section 35(2) of the Copyright Act 1968 'the author of a literary, dramatic, musical or artistic work is the owner of any copyright subsisting in the work'. By virtue of Section 32(1) copyright 'subsists in an original literary, dramatic, musical or artistic work that is unpublished' and of which the author was an Australian citizen, an Australian protected person or a person resident in Australia.

The Act, by Section 36(1) provides: 'Subject to this Act, the copyright in a literary, dramatic, musical or artistic work is infringed by a person who, not being the owner of the copyright and without the licence of the owner of the copyright, does in Australia, or authorises the doing in Australia of, any act comprised in the copyright'.

Section 31(1)(a)(i) provides that copyright includes the exclusive right to 'reproduce the work in a material form'. Thus, copyright is infringed by a person who, not being the owner of the copyright, reproduces or authorises the reproduction of a work, or of more than a reasonable part of the work, in a material form, unless the reproduction is a 'fair dealing' with the work 'for the purpose of research or study' as further defined in Sections 40 and 41 of the Act.

Section 51(2) provides that "Where a manuscript, or a copy, of a thesis or other similar literary work that has not been published is kept in a library of a university or other similar institution or in an archives, the copyright in the thesis or other work is not infringed by the making of a copy of the thesis or other work by or on behalf of the officer in charge of the library or archives if the copy is supplied to a person who satisfies an authorized officer of the library or archives that he requires the copy for the purpose of research or study'.

Modelling asset dynamics via an empirical
investigation of Australian Stock Exchange data

William Karel Bertram



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

School of Mathematics and Statistics
University of Sydney

May 2005

Abstract

This thesis is a study of the stochastic nature of financial markets and the modelling of asset dynamics. We review the main models for stock price dynamics. We compare the ability of the subordinated models to describe both daily and high frequency Australian stock price data. We look at the GARCH model and its ability to describe the correlation properties of the data.

We present an empirical study of high frequency Australian equity data. We examine the intraday behaviour of the Australian Stock Exchange (ASX), uncovering several properties including strong periodic trends, zero return enhancement, distributions with power law tails and correlations in various measures of the volatility. Power law exponents for the empirical distribution functions are estimated and compared with results from other studies. Using several different techniques we find the existence of long memory behaviour in the absolute log-return, volume and transaction frequency. We investigate the scaling behaviour of the data and show that high frequency Australian data exhibit a fractal scaling relationship.

We present a threshold model to describe the phenomena of zero return enhancement that is present in ASX data. This model depends on the occupation time distribution of an underlying market activity process. The occupation time distribution is calculated for several different models of activity. We fit the model to Australian data for small and large time scales and find that the model affords an excellent approximation of the distribution of stock returns.

The continuous time random walk (CTRW) is explored as a model for tick-by-tick price dynamics. We show that distribution of waiting times and survival function of ASX equities is well described by a non-Markovian CTRW where the memory kernel exhibits power law decay.

Acknowledgements

I would like to thank my supervisor, Dr. Peter Buchen, for his support and guidance. I also offer my thanks to Dr. Shelton Peiris and Dr. Hugh Luckock for their comments and assistance in developing this work. Finally I thank my friends and family for the support they have provided.

Data supplied by Securities Industry Research Centre of Asia-Pacific (SIRCA) on behalf of ASX.

Contents

1	Introduction	4
2	Models for price dynamics	13
2.1	Brownian motion	14
2.1.1	Wiener process and arithmetic Brownian motion	15
2.1.2	Geometric Brownian motion	17
2.2	Lévy-stable models	19
2.2.1	Stable laws and the Lévy process	20
2.3	Subordinated models	26
2.3.1	Properties of subordinated processes	28
2.3.2	The Mandelbrot model	30
2.3.3	The Clark lognormal model	31
2.3.4	The Student- t model	32
2.3.5	The variance gamma model	34
2.3.6	Comparison of models	36
2.4	Advanced subordinated models	39
2.4.1	The Poisson process	41
2.4.2	FATGBM	43
2.4.3	Multifractal model of asset returns	48
2.5	Jump-diffusion processes	49
2.5.1	The Merton model	50
2.5.2	The Cox-Ross model	51
2.6	Stochastic volatility	52

2.6.1	The Hull and White Model	53
2.6.2	The Ornstein-Uhlenbeck Model	54
2.6.3	The square root model	54
2.6.4	The 3/2 model	54
2.7	GARCH and generalised ARMA models	55
2.7.1	GARCH	55
2.7.2	Generalised ARMA models and misclassification	57
2.8	Summary	61
3	Tests and analysis of ASX equity data	64
3.1	Econophysics and empirical finance	65
3.2	ASX equity data	67
3.2.1	Overnight jumps and intraday noise	68
3.2.2	Intraday trend and master volatility measure	70
3.2.3	Business time	73
3.2.4	Zero return enhancement	80
3.3	Distribution properties	84
3.3.1	Parametric estimation	85
3.3.2	Empirical distribution function	90
3.3.3	Power law behaviour of equity data	95
3.3.4	The Hill estimator	100
3.4	Dependency properties	105
3.4.1	Autocorrelation and periodogram	107
3.4.2	Variance plots	110
3.4.3	Rescaled range analysis	112
3.4.4	Wavelet test for LRD	114
3.4.5	Detrended fluctuation analysis	116
3.5	Fractals and scaling	118
3.5.1	Scaling the empirical distribution	120
3.5.2	Scaling through the wavelet transform	127
3.5.3	Examining the fractal nature	129
3.6	Summary	132

4	New approaches for price modelling	135
4.1	Threshold Brownian motion	136
4.1.1	The model	138
4.1.2	Occupation time	140
4.1.3	Deterministic activity	143
4.1.4	Discrete stationary activity process	145
4.1.5	Brownian activity process	150
4.1.6	Ornstein-Uhlenbeck activity process	156
4.1.7	Results of fitting to ASX data	163
4.2	Montroll-Weiss continuous time random walk	166
4.2.1	CTRW theory	167
4.2.2	CTRW in finance	171
4.2.3	Empirical analysis of waiting times	173
4.3	Summary	181
5	Conclusion and closing remarks	184
A	Tables for Chapter 2	189
B	Tables for Chapter 3	197
C	Tables for Chapter 4	210
D	Programs for managing ASX data	213
E	Program for calculating the large deviation function	225

Chapter 1

Introduction

Over the past one hundred years there has been a growing demand for the mathematical modelling of financial markets. This period has seen a sharp increase in both the number and the complexity of the types of product available to investors. With each product comes a certain element of risk. Contracts such as: stocks, options, warrants and bonds can all be used via hedging, to increase the rate of return or decrease the risk of an investment. To achieve the appropriate outcomes it is necessary to measure accurately the risk associated with any given investment. For the purpose of risk management and portfolio theory, it is essential to know the frequency at which large standard deviation events occur. Therefore, it is of critical importance to understand the nature of the stochastic processes that drive financial markets.

At the end of the 19th century, Louis Bachelier developed the theory of random walks in the first attempt to describe the stochastic nature of price behaviour. His work was forgotten for almost sixty years until the famed economist Paul Samuelson rediscovered it in the 1950's. Bachelier is now credited as being the father of financial mathematics as well as having derived the theory of Brownian motion some 5 years before Albert Einstein.

The second half of the 20th century saw mathematicians begin to take an interest in financial and economic systems. Paul Samuelson revisited the random walk theory of financial markets. This led to the development of geometric Brownian motion

(GBM), the paradigm model for price dynamics. Early studies investigated the validity of GBM using historical data. Many of these studies reported the existence of fat tails in the log-return probability distribution, which could not be explained by geometric Brownian motion.

In 1963 Benoit Mandelbrot proposed the first fat tailed model for asset prices. His work is based on the stable Paretian hypothesis which states that the process of price change follows an α -stable Lévy process with $0 < \alpha < 2$. Lévy-stable distributions are stable attractors in probability space and are therefore the only limiting distributions for sums of independent identically distributed random variables. For $\alpha = 2$ the process is Brownian motion. However, if $0 < \alpha < 2$ the process has infinite variance. Although Lévy-stable processes are well defined mathematically, their property of infinite variance makes them extremely difficult to use in practice. Mandelbrot's work was revolutionary because it showed there were possible models other than GBM.

The work of Mandelbrot inspired the development of many more models. During the years following his initial contribution, the belief that the stock price exhibited a stochastic volatility, lead to the inception of the subordinated models. These models proposed that the price process could be described by a Brownian motion that evolved on a random time scale. Such models exhibit fat tails without the necessity of infinite variance.

The detection of correlations in financial time series led to the discrete time series models of ARCH and GARCH ((generalised) auto-regressive conditional heteroscedasticity). The two models remain extremely popular with practitioners due to their ease of implementation. However, the GARCH model is designed purely for fitting statistical properties of financial price data and offers little explanation about the process of price change.

When Black and Scholes introduced their option pricing model (Black and Scholes (1973)), which reduced the problem of pricing an option to one of solving a diffusion equation, they changed the nature of financial mathematics. Practitioners finally had a practical tool with which to price derivative contracts and for the next thirty years this type of problem would be of most importance to research.

The Black-Scholes model is based on the assumption of geometric Brownian motion. Over the years there have been several attempts to improve the model with extensions such as stochastic volatility models and jump-diffusion models. These amendments to the Black and Scholes model were made in order to ascertain how option prices behaved under alternative stock price models. The development of martingale methods in finance led to the influx of probability theorists into the field, bringing with them much of the theory and rigor of stochastic calculus. Although welcomed, such an addition has had the side effect of introducing a somewhat rigorous theory into what is still a poorly understood field.

Most of the models developed, from Lévy-stable processes to stochastic volatility, offer little insight into the process of price change. In many cases the models were proposed only to fit a particular distribution to the data. Furthermore there has been no general consensus as to which model best describes financial data. This comes down to a fundamental lack of understanding about the price change process itself. A good example of this is the controversy over the finite/infinite variance of financial time series.

The lack of agreement between models is largely owing to the fact that in the past there has been too little financial data available to study. The most important parts of a financial time series are the rare events. To observe enough of these large standard deviation events in order to draw statistical conclusions about them, one requires a time series with tens of thousands of entries. In the past, researchers could only work with daily or weekly data. With only ~ 250 trading days per year we would need 40 years worth of data to obtain a time series with 10,000 entries. There are obvious disadvantages in working with a time series over such a long sampling time. For instance, such series may be inhomogeneous due to changing market regulations or other factors. This would make it difficult to identify the properties of the stochastic process as distinct from those due to other factors.

With the advent of computerised stock exchanges, all information about the past state of the market is recorded electronically. Researchers now have the ability to access high-frequency financial data. Whereas in the past we used daily data, we now can look at 5 minute or 10 minute data and construct time series that are

hundreds of thousands of entries long. Such a time series may only span less than a decade in sampling time, thereby reducing the possibility that there is a significant change in the rules and regulations that govern the market.

High frequency data allows us to observe the market in greater detail than ever. With high frequency data we are able to determine which models best describe the properties of the data. More importantly we are able to further our understanding of the properties and phenomena exhibited by the stochastic processes of financial markets. With the coming of high frequency financial data there has been a movement by physicists to apply their theories and methods to financial markets. These physicists adopt the approach of seeking to understand the nature of the system. This approach is in stark contrast to the rigorous and theoretical emphasis of probability theory that has dominated financial mathematics for so long.

The emerging field of economic physics or econophysics, is focused on understanding the processes that drive financial markets by studying the empirical properties of high frequency data. The most investigated properties are undeniably those of the distribution, scaling behaviour and dependence structure. Recent interest has been in power law behaviour of financial distributions. Power laws first arose in a financial context due to the work of Pareto in the late 19th century. Mandelbrot introduced power laws into financial modelling with his stable Paretian hypothesis. Recent studies have found evidence of power law behaviour in financial markets, characterised by index values outside the realm of the stable distributions. Indeed, one of the main results of these studies is that there is now a general consensus among econophysicists that financial data are driven by a finite variance process.

Associated with power law behaviour is the property of scale invariance. The apparent scale invariance or fractal properties of stock prices have long been of interest to researchers. It is from such properties that we derive the analogy that all stock-price charts look the same. Understanding the nature of scaling behaviour is an important issue in financial research. Scaling behaviour can be exhibited by fractal or multifractal processes. For a fractal, the scaling relationship is linear, whilst for a multifractal it is non-linear. There is some evidence to suggest that multifractal behaviour is present in financial data. Several studies have found that

the distribution of log-returns seems to be scale invariant over small time scales, but which converge slowly to Gaussian over very large time scales.

There is also evidence of correlation structure in financial time series. The efficient market hypothesis leads to the price process being represented by a random walk of independent identically distributed random variables. Examination of the correlation of the log-return shows that the price increments are uncorrelated. However, new evidence has been unearthed of the possibility of a dependence structure in the stock's volatility. This existence of a dependence structure means that the random walk of independent identically distributed random variables cannot describe price change. Existence of such a property provides a reason as to why the central limit theorem does not ensure that the price process is Gaussian.

With the understanding gained by studying high frequency time series, the econophysicists can ultimately produce models that accurately represent the process of price change. Such models will provide investors with more accurate measures of the risks associated with any financial product.

Aims of this thesis

The aim of this thesis is to provide a comprehensive study of the modelling of stock price dynamics and gain an understanding of the stochastic processes that drive financial markets. Central to this aim is the proposal of models that address specific phenomena observed in the data. It is through these models that we learn how financial markets behave and how the different processes interact with each other. This provides us with the knowledge and ability to develop more sophisticated and more realistic models.

To this end we will present a review of the major developments, research directions and trends in asset modelling from the past century. As part of this review we will examine the theory of previously proposed models and compare the abilities of those models to describe the distribution behaviour of Australian Stock Exchange daily price data. In order to examine better the behaviour of the processes involved in financial markets we will use high frequency financial data to conduct an empirical study of Australian equity data. The purpose of this study is to provide information

and knowledge about the fundamental stochastic nature of the Australian market. Further, we will use high frequency data to examine the performance of the previous models for price change, focusing especially on the tail behaviour of the models. Finally, we will examine new models to describe high frequency price dynamics. These new models provide explanations as to why various phenomena are observed in market data and we will show that these new models are able to describe phenomena that previous models could not.

The organisation of the thesis is as follows. In Chapter 2 we review the history of models for price dynamics. Starting with Bachelier we follow the developments that lead to geometric Brownian motion. We examine the stable Paretian model of Mandelbrot and the theory of Lévy-stable distributions. We look at the theory of subordinated processes and investigate four of the main models. The performance of each model is then examined using daily price data from the ASX. We consider several more recent models that extend the idea of subordination to include a correlation structure. The asset models with stochastic volatility and jump-diffusion are reviewed. We also look at the GARCH model and fit GARCH(1,1) to ASX daily price data. The final part of the chapter investigates a misclassification problem that can arise when forecasting on time series of varying frequency content.

Chapter 3 focuses on the real world behaviour of assets through the use of high frequency financial data. The fundamental rules of the ASX will be shown to influence the price dynamics of the securities that trade on it. Markets such as the ASX allow for price development during non-trading periods, such as after hours trading and dual-listed stocks. A new approach for working with securities that trade on exchanges similar to the ASX will be introduced. This approach amounts to representing the stock returns as two separate stochastic processes, a discontinuous overnight return process and a continuous intraday return process. We also propose a new measure for the market activity and use it to normalise the average market activity levels during trading hours. We examine the distribution properties of the log-return, volume, and transaction frequency. We find evidence of a disproportionate number of zero returns and estimate the tail behaviour empirically and by using Hill estimates. We detect the existence of long memory in the absolute log-return,

volume and transaction frequency. These results are compared to results found in other studies and indicate that the behaviour of Australian equities is significantly different from the reported behaviour of other financial markets. We examine the scaling behaviour of the ASX data and show that for high frequency time intervals the data is adequately described by a fractal scaling relationship.

In Chapter 4 we present a new two factor model for the behaviour of price dynamics. This model offers a possible explanation of the origin of the large number of zero returns found in equity data from the Australian Stock Exchange. Our model is able to explain phenomenologically the disproportionate number of zero returns and show how it is related to the activity level of the market place. With the model we are able to demonstrate that zero return enhancement over a small time interval affects the stock return over a large time interval. We examine the model under four different representations of market activity and fit the model to ASX log-return data on both large and small time scales. We investigate the continuous time random walk (CTRW) as a model for high frequency stock prices. We review the theory of the CTRW and show that this model is able to reproduce the waiting time distribution for ASX stocks under an appropriate choice of memory kernel.

Chapter 5 concludes this thesis with a summary of the major results and findings of this thesis.

Original contributions made in this thesis

This thesis contains several original contributions to the subjects of econophysics and financial mathematics. In Chapter 2 we provide a comparison of subordinated models on both daily data and high frequency data. We find that the model of Praetz (1972) is the best fitting for daily data while the variance gamma model of Madan and Seneta (1990) is the best performing subordinated model on high frequency data. We also investigate a misclassification problem that has strong implications for forecasting on time series. We show that the GAR model of Peiris (2004) improves forecasting performance.

Chapter 3 makes a contribution to the subject through a comprehensive empirical analysis of high frequency data. It is important for studies such as this to be

carried out as financial properties can be dependent on the individual microstructure of an exchange as well as the geographic location and social behaviour of market participants. We present the intraday trend of the absolute log-return, volume and transaction frequency. We introduce a new measure for the volatility constructed through principal component analysis. We find evidence of power law behaviour in financial data whose index α differs to that found in several other studies. This finding has implications for a recent model for financial power law behaviour by Gabaix et.al. (2003a). We detect and measure long range dependence in ASX time series. We find evidence that shows ASX high frequency data to possess fractal properties.

In Chapter 4 the original contribution takes the form of a new model for asset dynamics. This model represents the first attempt to address the phenomenon of zero return enhancement. We make a further contribution by investigating the properties of the CTRW model using high frequency ASX data. We find evidence of a correlation between the log-return and the waiting time. We also show that the empirical survival function displays behaviour consistent with the Mittag-Leffler function.

Published work contained in this thesis

The original work contained in this thesis has given rise to the following publications,

- W.K. Bertram, An empirical investigation of Australian Stock Exchange data, *Physica A*, 341 (2004), 533-546.
- W.K. Bertram, A threshold model for Australian equities, *Physica A*, 346 (2005), 561-577.
- W.K. Bertram and M.S. Peiris, An example of a misclassification problem applied to Australian equity data, *Submitted to Economics Letters*.

A note on presentation

Due to the empirical nature of the subject, this thesis contains many figures and tables that provide the results of various analysis. In order to improve readability,

the figures are incorporated into the body of the text, while the tables are relegated to the appendices. Appendix A contains the tables that are referenced in Chapter 2, Appendix B contains those of Chapter 3 and Appendix C contains those of Chapter 4.

Chapter 2

Models for price dynamics

Mathematical modelling of financial markets extends back for over a century. There have been many contributions to this subject over the years from a broad spectrum of the sciences. This chapter details the history of asset modelling over the past hundred years and provides a review of the main models and their properties. Beginning with Bachelier's original model we progress through the paradigm model of geometric Brownian motion. The continuing use of GBM is largely due to the huge success of the Black-Scholes option pricing model coupled with GBM's ability to achieve analytic results. We examine Mandelbrot's many contributions to financial modelling starting with the stable Paretian hypothesis and the ongoing work on multifractals. The important class of subordinated asset models is reviewed and a comparison of the various models is undertaken using daily stock return data from the Australian Stock Exchange. Building on the earlier work of subordination we explore some recently developed models that extend the method of subordination to a more general form of compounding process. This development allows for the incorporation of a correlation structure into price modelling. We review the jump-diffusion model of Merton and the alternate asset models proposed by Cox and Ross. This is followed by a survey of the main stochastic volatility models. We also review the conditional distribution discrete time series models of ARCH and GARCH. These discrete time series models have become an important tool for financial practitioners. In the final section of the chapter we present the work of Bertram and Peiris (2004), where we

examine a misclassification problem that can arise when modelling data of varying frequency content.

--- Brownian motion

On the 29th of March, 1900, the French mathematician Louis Bachelier presented his doctoral thesis at the Sorbonne in Paris. He defended his work before a board consisting of some of the most eminent mathematicians of the late 19th and early 20th century: Joseph Boussinesq, Paul Appell and Henri Poincaré. His thesis, which was favorably remarked upon by Poincaré, contained some of the most ground breaking work in what are now the fields of stochastic calculus and financial mathematics.

Bachelier's thesis, titled "*Théorie de la Speculation*" contains a theoretical study of the random nature of the stock market. His work marks the beginning of the theory of stochastic processes. Contained in his work is a derivation of what we today know as the Chapman-Kolmogorov equation for continuous stochastic processes, the diffusion equation that determines the probability of a price change. Bachelier went on to derive the Wiener process and recognise that its probability density function is a solution of the partial differential equation for heat diffusion, something that would be redone by Einstein five years later in his classic paper "*Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen*". Bachelier continued on to solve the first passage time problem for the Wiener process and constructed the price for a barrier option. Despite the importance of Bachelier's work, his discoveries were not appreciated until decades later when the work of Samuelson (1955), Krueger (1956), Osbourne (1959) and Cootner (1964) reported the extent of Bachelier's contributions. Bachelier's work is quite heuristic in approach and was scorned at the time by mathematicians who were in the process of transforming probability theory into a rigorous discipline. Most of the theory contained in his thesis was later rediscovered independently by the likes of Lévy, Kolmogorov, Borel, Wiener, Khinchine, Markov and Feller¹. Today with the enormous importance placed on financial contracts and

¹William Feller had originally called the Wiener process the Bachelier-Wiener Process.

the expanding field of financial mathematics, Bachelier's work is finally appreciated in its proper context and is rightly recognised as the pioneering work in this field.

2.1.1 Wiener process and arithmetic Brownian motion

The Wiener process, W_t was proposed by Wiener (1923) as a mathematical description of Brownian motion. It is a Gaussian, continuous time stochastic process with stationary, independent increments for which,

$$W_0 = 0; \quad (2.1)$$

$$\mathbb{E}[W_t] = 0; \quad (2.2)$$

$$\mathbb{V}[W_t] = t; \quad (2.3)$$

$$\text{cov}[W_t, W_s] = \min(t, s); \quad \text{for all } s, t \in \mathbb{R}. \quad (2.4)$$

From the above properties the Wiener process is distributed as a normal random variable with mean $\mu = 0$ and variance $\sigma^2 = t$,

$$W_t \sim N(0, t). \quad (2.5)$$

While the increments are stationary, the Wiener process is non-stationary in the strict sense. That is, its joint distributions are time dependent and not invariant under a time shift. Furthermore, the Wiener process is space homogeneous,

$$W_t^x = x + W(t) \quad (2.6)$$

where W_t^x is the Wiener process started at $W(0) = x$. The Wiener process is a martingale,

$$\mathbb{E}[W_{t+s} | \mathcal{F}_t] = W_t; \quad s > 0 \quad (2.7)$$

where \mathcal{F}_t is the filtration (or history of the process) up to time t . Also the Wiener process possesses the Markov property,

$$\mathbb{P}[W_{t+s} \leq y | \mathcal{F}_t] = \mathbb{P}[W_{t+s} \leq y | W_t]. \quad (2.8)$$

This property relays the fact that the history of the Wiener process is entirely specified by its state at the previous time step. The Wiener process is self-similar,

i.e. scale invariant,

$$W_{at} \stackrel{d}{=} \sqrt{a}W_t \quad (2.9)$$

where $\stackrel{d}{=}$ denotes equality in distribution. Thus changing the time scale on which the process operates amounts to rescaling the amplitude of the Wiener process. Turning our attention to the increments of the Wiener process, $\Delta W_s = W_{s+\Delta t} - W_s$, we find that they possess the following properties,

$$\mathbb{E}[\Delta W_t] = 0 \quad (2.10)$$

$$\mathbb{V}[\Delta W_t] = \Delta t \quad (2.11)$$

$$\text{corr}[\Delta W_t, \Delta W_s] = \begin{cases} 1; & \text{if } s = t \\ 0; & \text{if } s \neq t \end{cases} \quad (2.12)$$

for $s \notin (t + \Delta t)$, $t \notin (s + \Delta t)$. Hence the increment process is distributed as independent normal random variables with mean $\mu = 0$ and variance $\sigma^2 = \Delta t$,

$$\Delta W_t \sim N(0, \Delta t). \quad (2.13)$$

Unlike the Wiener process, the increments ΔW_t are strictly stationary,

$$\Delta W_t \stackrel{d}{=} \Delta W_s; \quad \text{for all } t, s \quad (2.14)$$

as the increment's distribution is only dependent on Δt , rather than on t itself.

Using the framework of modern stochastic calculus we recognise Bachelier's stock price model as being that of arithmetic Brownian motion. It can be expressed using the stochastic differential equation representation,

$$dP_t = \mu dt + \sigma dW_t \quad (2.15)$$

where μ and σ^2 are the instantaneous drift and variance of the process P_t . This stochastic differential equation has solution,

$$P_t = P_0 + \mu t + \sigma W_t \quad (2.16)$$

where P_0 is the price of the stock at time $t = 0$. Arithmetic Brownian motion inherits the previously described features of the Wiener process, namely stationary

independent increments and the Markov property. Under the arithmetic Brownian motion model, the stock price return, $R_t = P_{t+\Delta t} - P_t$ is the quantity of most interest. This value describes the random change in price over a time interval $[t, t + \Delta t]$. The return is normally distributed,

$$\mathbf{P}[R_t \leq x] = \mathcal{N}\left(\frac{x - \mu\Delta t}{\sigma\sqrt{\Delta t}}\right) \quad (2.17)$$

where $\mathcal{N}(x)$ is the normal distribution function:

$$\mathcal{N}(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{u^2}{2}} du. \quad (2.18)$$

A fundamental property of stock prices is that while they are unbounded above, they are bounded below by zero, *i.e.* $P_t \geq 0$. It is important to note that Bachelier's model allows for negative stock price. According to Cootner (1964), Bachelier recognised the problem posed by the asymmetry of the distribution of stock prices but ignored it. By dealing with *rentes* (government bonds) Bachelier avoided complications posed by the fact that stock prices have a lower bound at zero because over a short period, it is unlikely to find a price change which is influenced in any major way by the barrier at zero. More than fifty years later Kruizenga (1956), motivated by Paul Samuelson, used similar methods to Bachelier to price put and call options.

2.1.2 Geometric Brownian motion

To address the problem that allowed for negative prices, Samuelson (1955), Osborne (1959) and Alexander (1961) proposed that rather than modelling the return by a normal distribution, such a distribution should model the logarithmic return. This idea followed from the rate of return proposition, reiterated in Samuelson (1965), which states that it is the proportional return that interests investors rather than the absolute return, and that every dollar of market value is subject to the same percentage fluctuations per unit time regardless of the absolute value of the stock. This model became known as geometric Brownian motion and is described by the stochastic differential equation,

$$dP_t = \mu P_t dt + \sigma P_t dW_t. \quad (2.19)$$

Through the stochastic calculus developed by Itô (1951), the solution to this equation is found to be

$$P_t = P_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma W_t}. \quad (2.20)$$

For geometric Brownian motion, we find that the log-price follows an arithmetic Brownian motion. The quantity of interest is now the log-return over time interval length Δt , defined as,

$$L_t = \log P_{t+\Delta t} - \log P_t \quad (2.21)$$

$$= (\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma[W_{t+\Delta t} - W_t]. \quad (2.22)$$

Since $W_{t+\Delta t} - W_t \sim N(0, \Delta t)$ we see that the log-return is indeed distributed as a normal random variable,

$$\mathbf{P}[L_t \leq x] = \mathcal{N}\left(\frac{x - (\mu - \frac{1}{2}\sigma^2)\Delta t}{\sigma\sqrt{\Delta t}}\right). \quad (2.23)$$

The implications of using such a model for stock price returns are that the market is required to follow several strict assumptions. The main assumption is that the log-returns are independent identically distributed normal random variables. This assumption has been the focus of much study over the past fifty years. During the early 1960s, several studies were undertaken by authors including Osbourne (1959), Alexander (1961), Kendall (1953), Moore (1962) and the collected works in Cootner (1964) to investigate whether this hypothesis was true. Many early studies such as Mandelbrot (1963), Fama (1963, 1965) were quick to reject this hypothesis and proposed improved methods to model asset returns better.

In 1973, seventy three years after Bachelier first attempted to price an option contract, Fisher Black and Myron Scholes introduced their celebrated option pricing model (Black and Scholes (1973)). The Black-Scholes model was based on the assumption that the stock price follows a geometric Brownian motion. For the first time, this framework obtained the price of options independent of investor risk profile. This was a sensation and had an immediate effect on the world of finance. In the twenty years following the inception of the Black-Scholes pricing model, option pricing became the dominant form of financial mathematics while asset modelling, though still actively researched, took somewhat of a back seat. Following the market

crash of 1987 and several spectacular incidents in the 1990's, the fallibility of Black-Scholes through its use of geometric Brownian motion was realised and interest was once again focused on developing models that better reflect the reality of stock prices.

2.2 Lévy-stable models

The assumption of the Gaussian nature of stock returns is derived from the belief that price changes are independent identically distributed random variables with finite variance. If the transaction frequency is fairly evenly spaced through time then the central limit theorem infers that the price changes will be normally distributed. Empirical studies of the log-return distribution carried out in the early 1950s and 1960s seemed to reinforce this belief, however most of these studies indicated that the tails of the empirical distributions might be fatter than those required by the Gaussian assumption.

In 1963 Benoit Mandelbrot proposed a new approach to the random walk theory of finance with the introduction of the family of Lévy-stable distributions. Although Mandelbrot's work was seen as revolutionary at the time, the motivation for his work can be traced back to the work of Vilfredo Pareto in 1896. Pareto (1896) observed that in a given population, the distribution of personal incomes, I , above a given level, x , behaves like $\mathbf{P}[I > x] \sim x^{-\alpha}$ with $\alpha > 0$. One of the main properties of the stable distributions is that their distribution tails display this type of power law behaviour. Mandelbrot (1959, 1960) stated that the stable laws provide the appropriate model for income distributions. Subsequently, Mandelbrot (1963) comprised an empirical study of commodity spot prices arguing a case for the stable Paretian hypothesis over that of geometric Brownian motion. This work was met with mixed reactions. Firstly, as discussed by Cootner (1964), the value of commodity prices is strongly related to the volume of inventories held. Thus one would not expect commodity spot prices to follow a geometric Brownian motion in the first place. Secondly, the Lévy distributions introduce the concept of infinite variance and the meaninglessness of correlations. The high probability of large

changes in the stock return makes statistical techniques based on finite variance inapplicable. This property means that most of the current tools used to deal with stochastic processes are made redundant, a prospect not very appealing for practitioners. Like Bachelier before him, Mandelbrot too takes a somewhat heuristic approach and relies heavily on qualitative measures. In Mandelbrot (1963) he argued quite strongly in favor of the stable hypothesis, and as Cootner (1965) noted, the determination and passion with which he argued

“... lend it a messianic tone which makes me just a bit uncomfortable.”

Nevertheless Mandelbrot’s work forced researchers to face the fact that the Gaussian random walk does not represent the reality of stock prices and that alternative approaches are possible, even desirable.

2.2.1 Stable laws and the Lévy process

The family of Lévy-stable distributions was first described by Paul Lévy in “*Calcul des Probabilités*” (1925). Lévy processes are stochastic processes that follow a generalised central limit theorem.

Suppose X_1, X_2, \dots, X_n are independent identically distributed random variables with common distribution function $F(x)$. Then the distribution $F(x)$ is said to be stable if for any n ,

$$a_n(X_1 + X_2 + \dots + X_n) + b_n \stackrel{d}{=} X_1 \quad (2.24)$$

with $a_n > 0$ and $b_n \in \mathbb{R}$. A stable random variable has the property that the functional form of its density is preserved under addition. Consider the partial sums of a sequence of stable random variables,

$$S_n = \sum_{i=1}^n X_i \quad (2.25)$$

with density function $f(S_n)$. The functional form of $f(S_n)$ will be the same as that of $f(X_i)$. It is sometimes more intuitive to examine the property of stability in terms of the characteristic function of a random variable. The characteristic function is

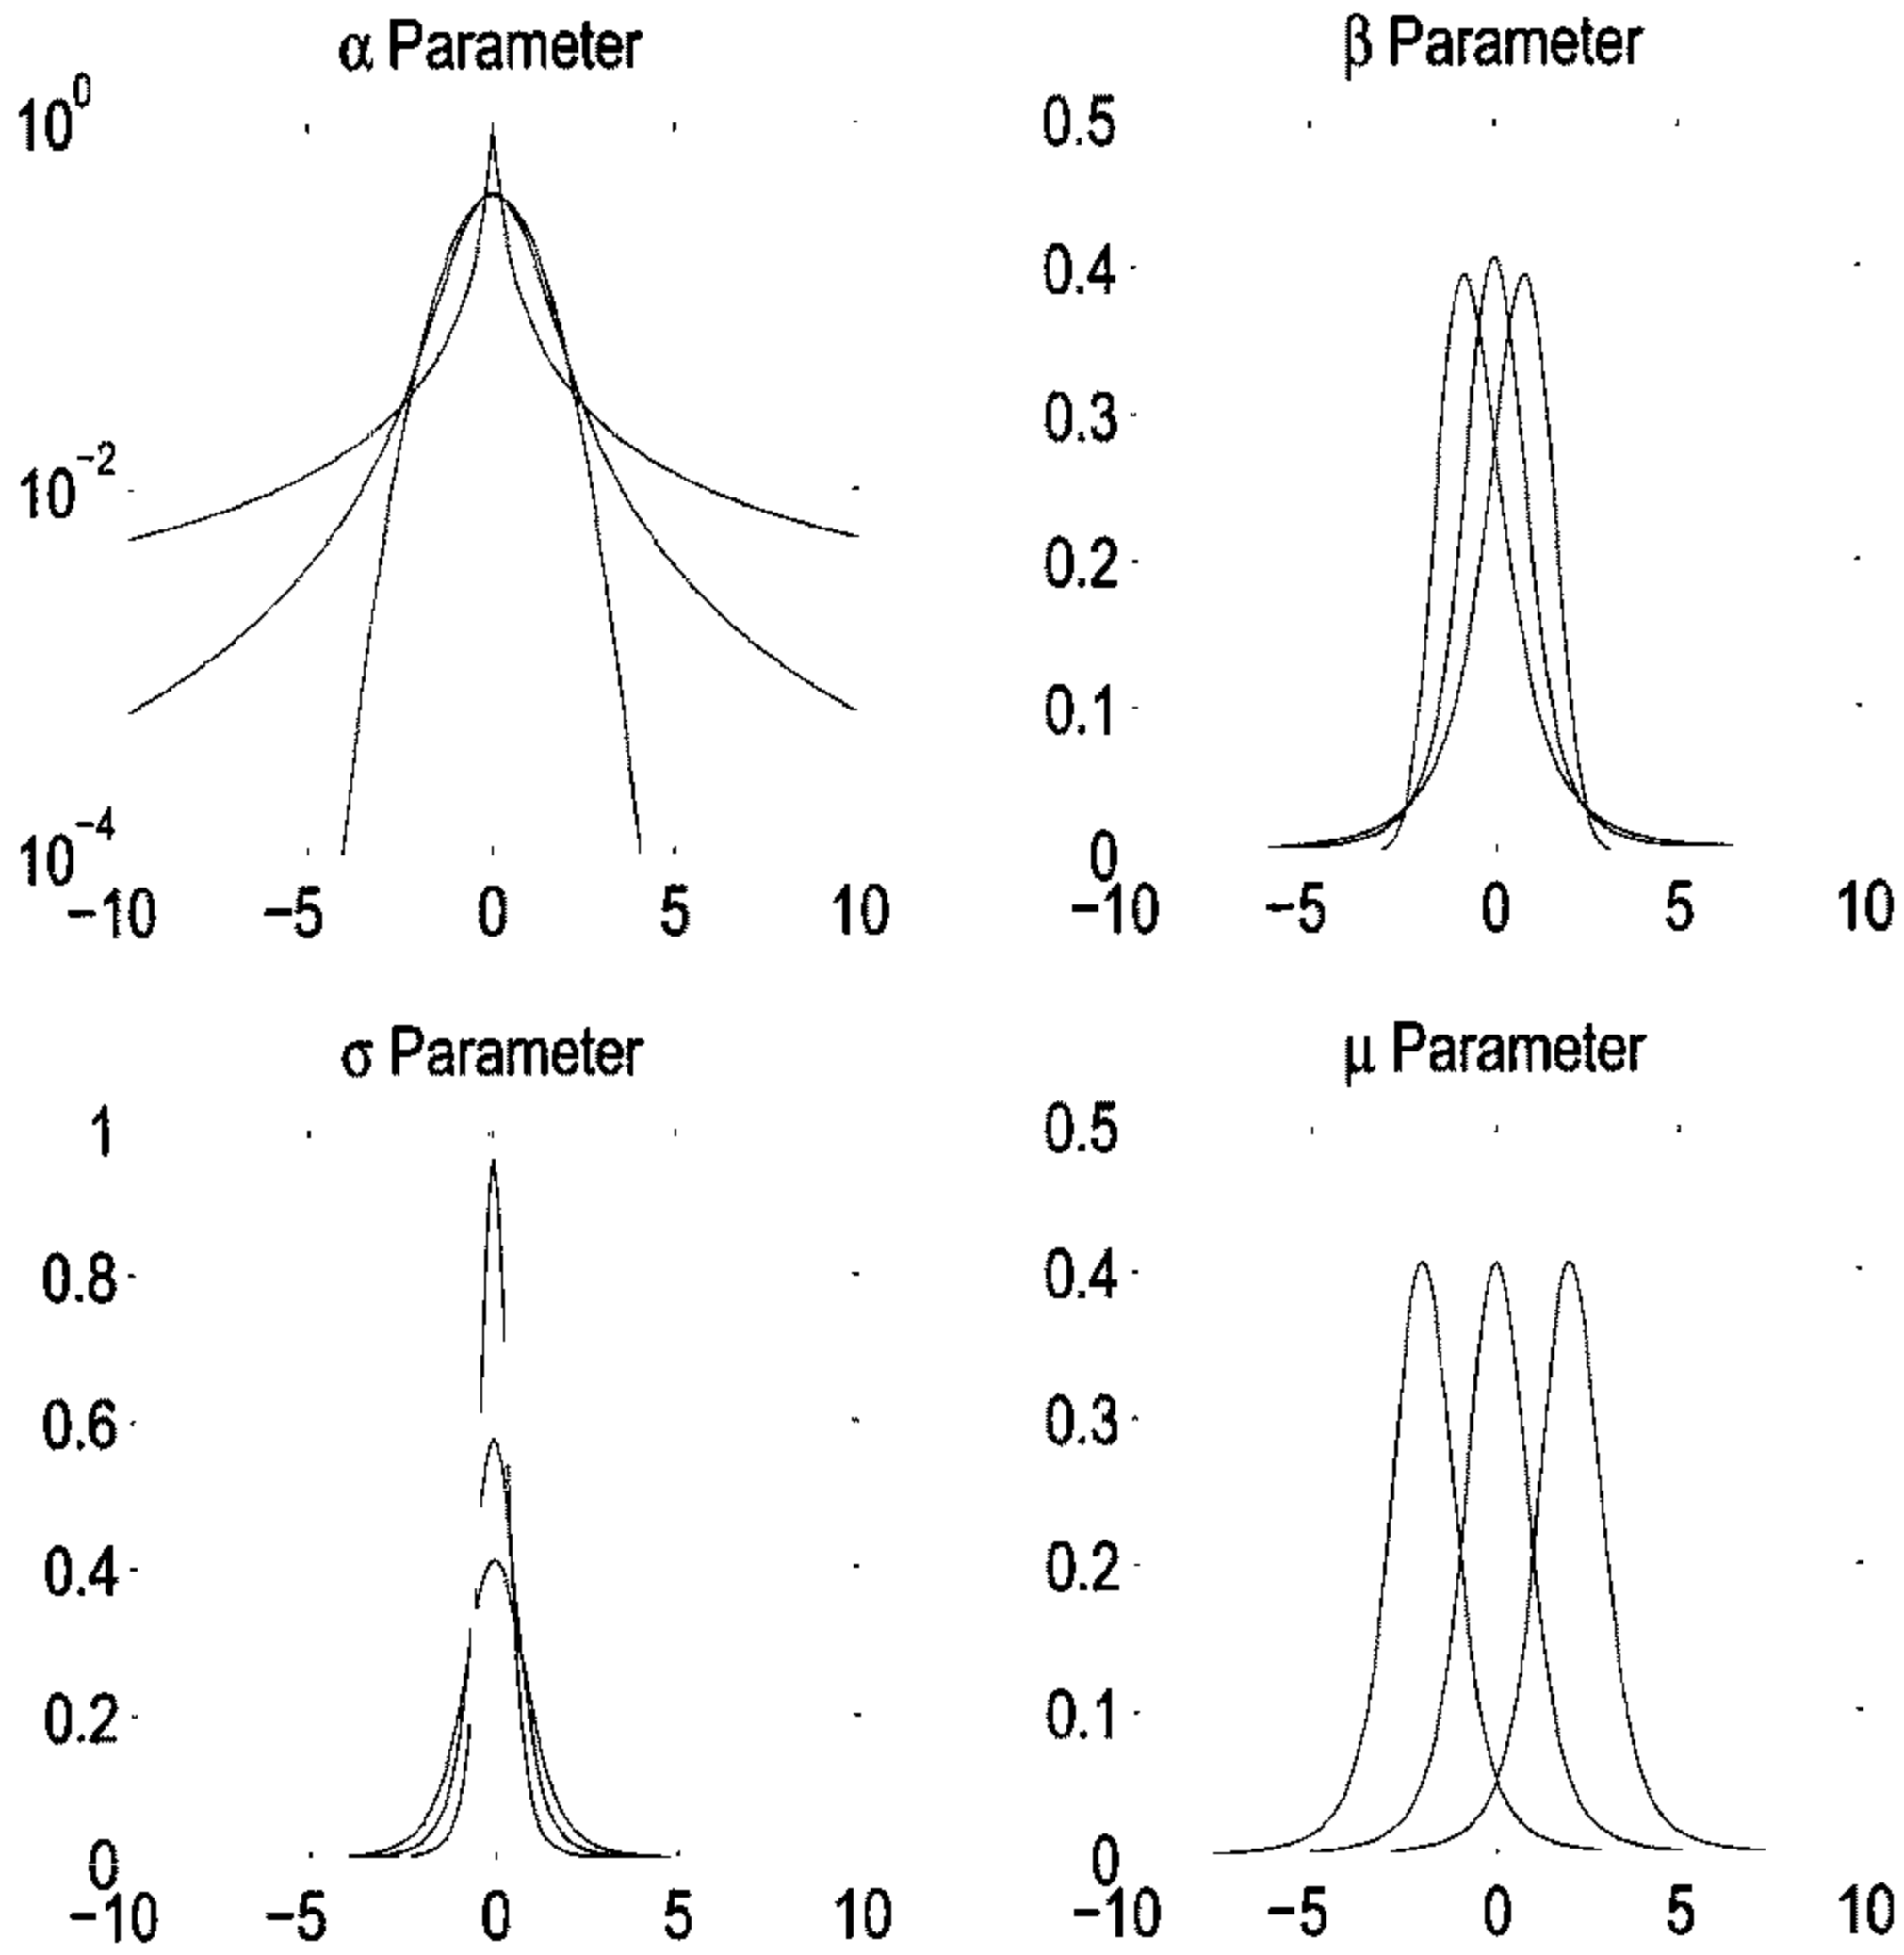


Figure 2.1: This figure displays the Lévy-stable density functions for different parameter values. The plots in this figure show the different effects that can be achieved by altering the parameters: α changes the tail fatness, β the skewness, σ the shape (width) and μ the position.

defined as the Fourier transform of the distribution function,

$$\phi(\theta) = \int_{-\infty}^{\infty} f(x)e^{i\theta x} dx. \quad (2.26)$$

For independent random variables, x_1, x_2 the distribution of the sum

$$S_2 = x_1 + x_2 \quad (2.27)$$

is given by the convolution of the densities $f(x_1)$ and $f(x_2)$,

$$f(S_2) = f(x_1) \otimes f(x_2). \quad (2.28)$$

Thus for the partial sums defined by Eq. (2.25), we can write the density function $f(S_n)$ as

$$f(S_n) = f(X_1) \otimes f(X_2) \otimes \dots \otimes f(X_n) \quad (2.29)$$

$$= f(X_1) \otimes f(X_1) \otimes \dots \otimes f(X_1). \quad (2.30)$$

Distributions that satisfy this equation are often referred to as being closed under convolution. Taking a Fourier transform of Eq. (2.30) we use the convolution theorem to write the equation in terms of a product of characteristic functions,

$$\phi_{S_n}(\theta) = \phi(\theta)\phi(\theta)\dots\phi(\theta) \quad (2.31)$$

$$= (\phi(\theta))^n. \quad (2.32)$$

The work of Lévy (1925) and Khintchine (1936) determined that the most general form of characteristic function that is preserved under Eq. (2.32) is given by

$$\phi(\theta) = \begin{cases} \exp[i\mu\theta - \sigma^\alpha|\theta|^\alpha[1 - i\beta\frac{\theta}{|\theta|}\tan(\frac{\pi\alpha}{2})]]; & \alpha \neq 1 \\ \exp[i\mu\theta - \sigma|\theta|[1 + i\beta\frac{\theta}{|\theta|}\frac{2}{\pi}\ln|\theta|]]; & \alpha = 1, \end{cases} \quad (2.33)$$

where $0 < \alpha \leq 2$, $\sigma > 0$, $\mu \in \mathbb{R}$, and $-1 < \beta < 1$. The family of distributions with characteristic function given by Eq. (2.33) are known as the Lévy-stable distributions. These distributions are the only possible limiting distributions for sums of independent identically distributed random variables. The Lévy-stable distributions are a four parameter family, usually written as $S_\alpha = S_\alpha(\sigma, \beta, \mu)$. Fig. 2.1 displays

the effect that the different parameters have on the form of the distribution. When $\mu = 0$, $\beta = 0$ we have the set of zero mean, symmetric stable distributions with characteristic function,

$$\phi(\theta) = e^{-\sigma|\theta|^\alpha}. \quad (2.34)$$

In the case where $\alpha = 2$ the characteristic function reduces to that of the normal distribution.

The Lévy-stable density function can be expressed as the inverse Fourier transform of the characteristic function,

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\sigma|\theta|^\alpha} e^{-i\theta x} d\theta. \quad (2.35)$$

The only known Lévy distributions that possess analytic forms for Eq. (2.35) are: the Lévy-Smirnov ($\alpha = 1/2$, $\beta = 1$), the Lorentzian ($\alpha = 1$, $\beta = 0$) and the Gaussian ($\alpha = 2$) distributions. When $\sigma = 1$ there exists a series expansion for large $|x|$ (Mantegna and Stanley (2000)),

$$f(x) = -\frac{1}{\pi} \sum_{k=1}^n \frac{(-1)^k \Gamma(\alpha k + 1)}{k! |x|^{\alpha k + 1}} \sin \left[\frac{k\pi\alpha}{2} \right] + R(|x|) \quad |x| \rightarrow \infty \quad (2.36)$$

where the remainder term is,

$$R(|x|) = O(|x|^{-\alpha(n+1)-1}). \quad (2.37)$$

Thus the asymptotic approximation of the Lévy-stable density for large $|x|$ indicates that the distribution tails display power law behaviour,

$$f(|x|) \sim \frac{\Gamma(1 + \alpha) \sin(\pi\alpha/2)}{\pi |x|^{1+\alpha}} \sim |x|^{-(1+\alpha)}. \quad (2.38)$$

For $0 < \alpha < 2$ the distribution tails will always be fatter than those of the normal distribution.

The α -stable Lévy process is the stable analogue of the Wiener process. In the same way that the normal distribution satisfies the Chapman-Kolmogorov partial differential equation (CKE), the Lévy-stable distributions also satisfy the CKE, but only when the condition of finite variance is relaxed. Unlike the Wiener process and

Brownian motion, the Lévy process is discontinuous due to the presence of large jumps with infinite variance.

The α -stable Lévy process (with $\beta=0$, $\mu = 0$) is symmetric and has independent increments,

$$X_t - X_s \sim S_\alpha(|t - s|^{\frac{1}{\alpha}}, 0, 0). \quad (2.39)$$

The process possesses stationary increments,

$$X_{t+s} - X_s \stackrel{d}{=} X_t - X_0 \quad (2.40)$$

and the Lévy process itself is self-similar with scaling parameter $H = \frac{1}{\alpha}$,

$$X_{ct} \stackrel{d}{=} c^{\frac{1}{\alpha}} X_t. \quad (2.41)$$

In all cases when $\alpha \rightarrow 2$ these results reduce to those of Brownian motion.

Mandelbrot's stable Paretian hypothesis implies that the unconditional distribution of log-returns is distributed as a Lévy-stable symmetric random variable with infinite second moments,

$$\mathbf{P}[L_t \leq x] \sim S_\alpha(\sigma t^{\frac{1}{\alpha}}, 0, \mu) \quad (2.42)$$

where $L_t = \log P_{t+\tau} - \log P_t$, $0 < \alpha < 2$. Much work has been done investigating whether stock returns have infinite second moments and this remains a contentious issue between researchers today.

Using historical records for the last traded price we can fit the four parameter Lévy distribution to ASX equity data to observe the description offered by the stable Paretian hypothesis. Note that Mandelbrot specified a three parameter model, whereas here we have chosen to include the skewness parameter, β in order to demonstrate the descriptive power of the model. Fig. 2.2 shows the empirical density function for the log-returns calculated for several stocks along with estimated Lévy-stable densities. Table A.1 displays the parameter values gained by fitting the Lévy densities to the data set. We find that for daily returns the index parameter lies in the range $\alpha \sim 1.7 \pm 0.1$. The skewness parameter β varies from stock to stock, whereas the shape parameter σ are all the same and the mean μ is almost zero for all stocks. In this case the low frequency nature of the data means that

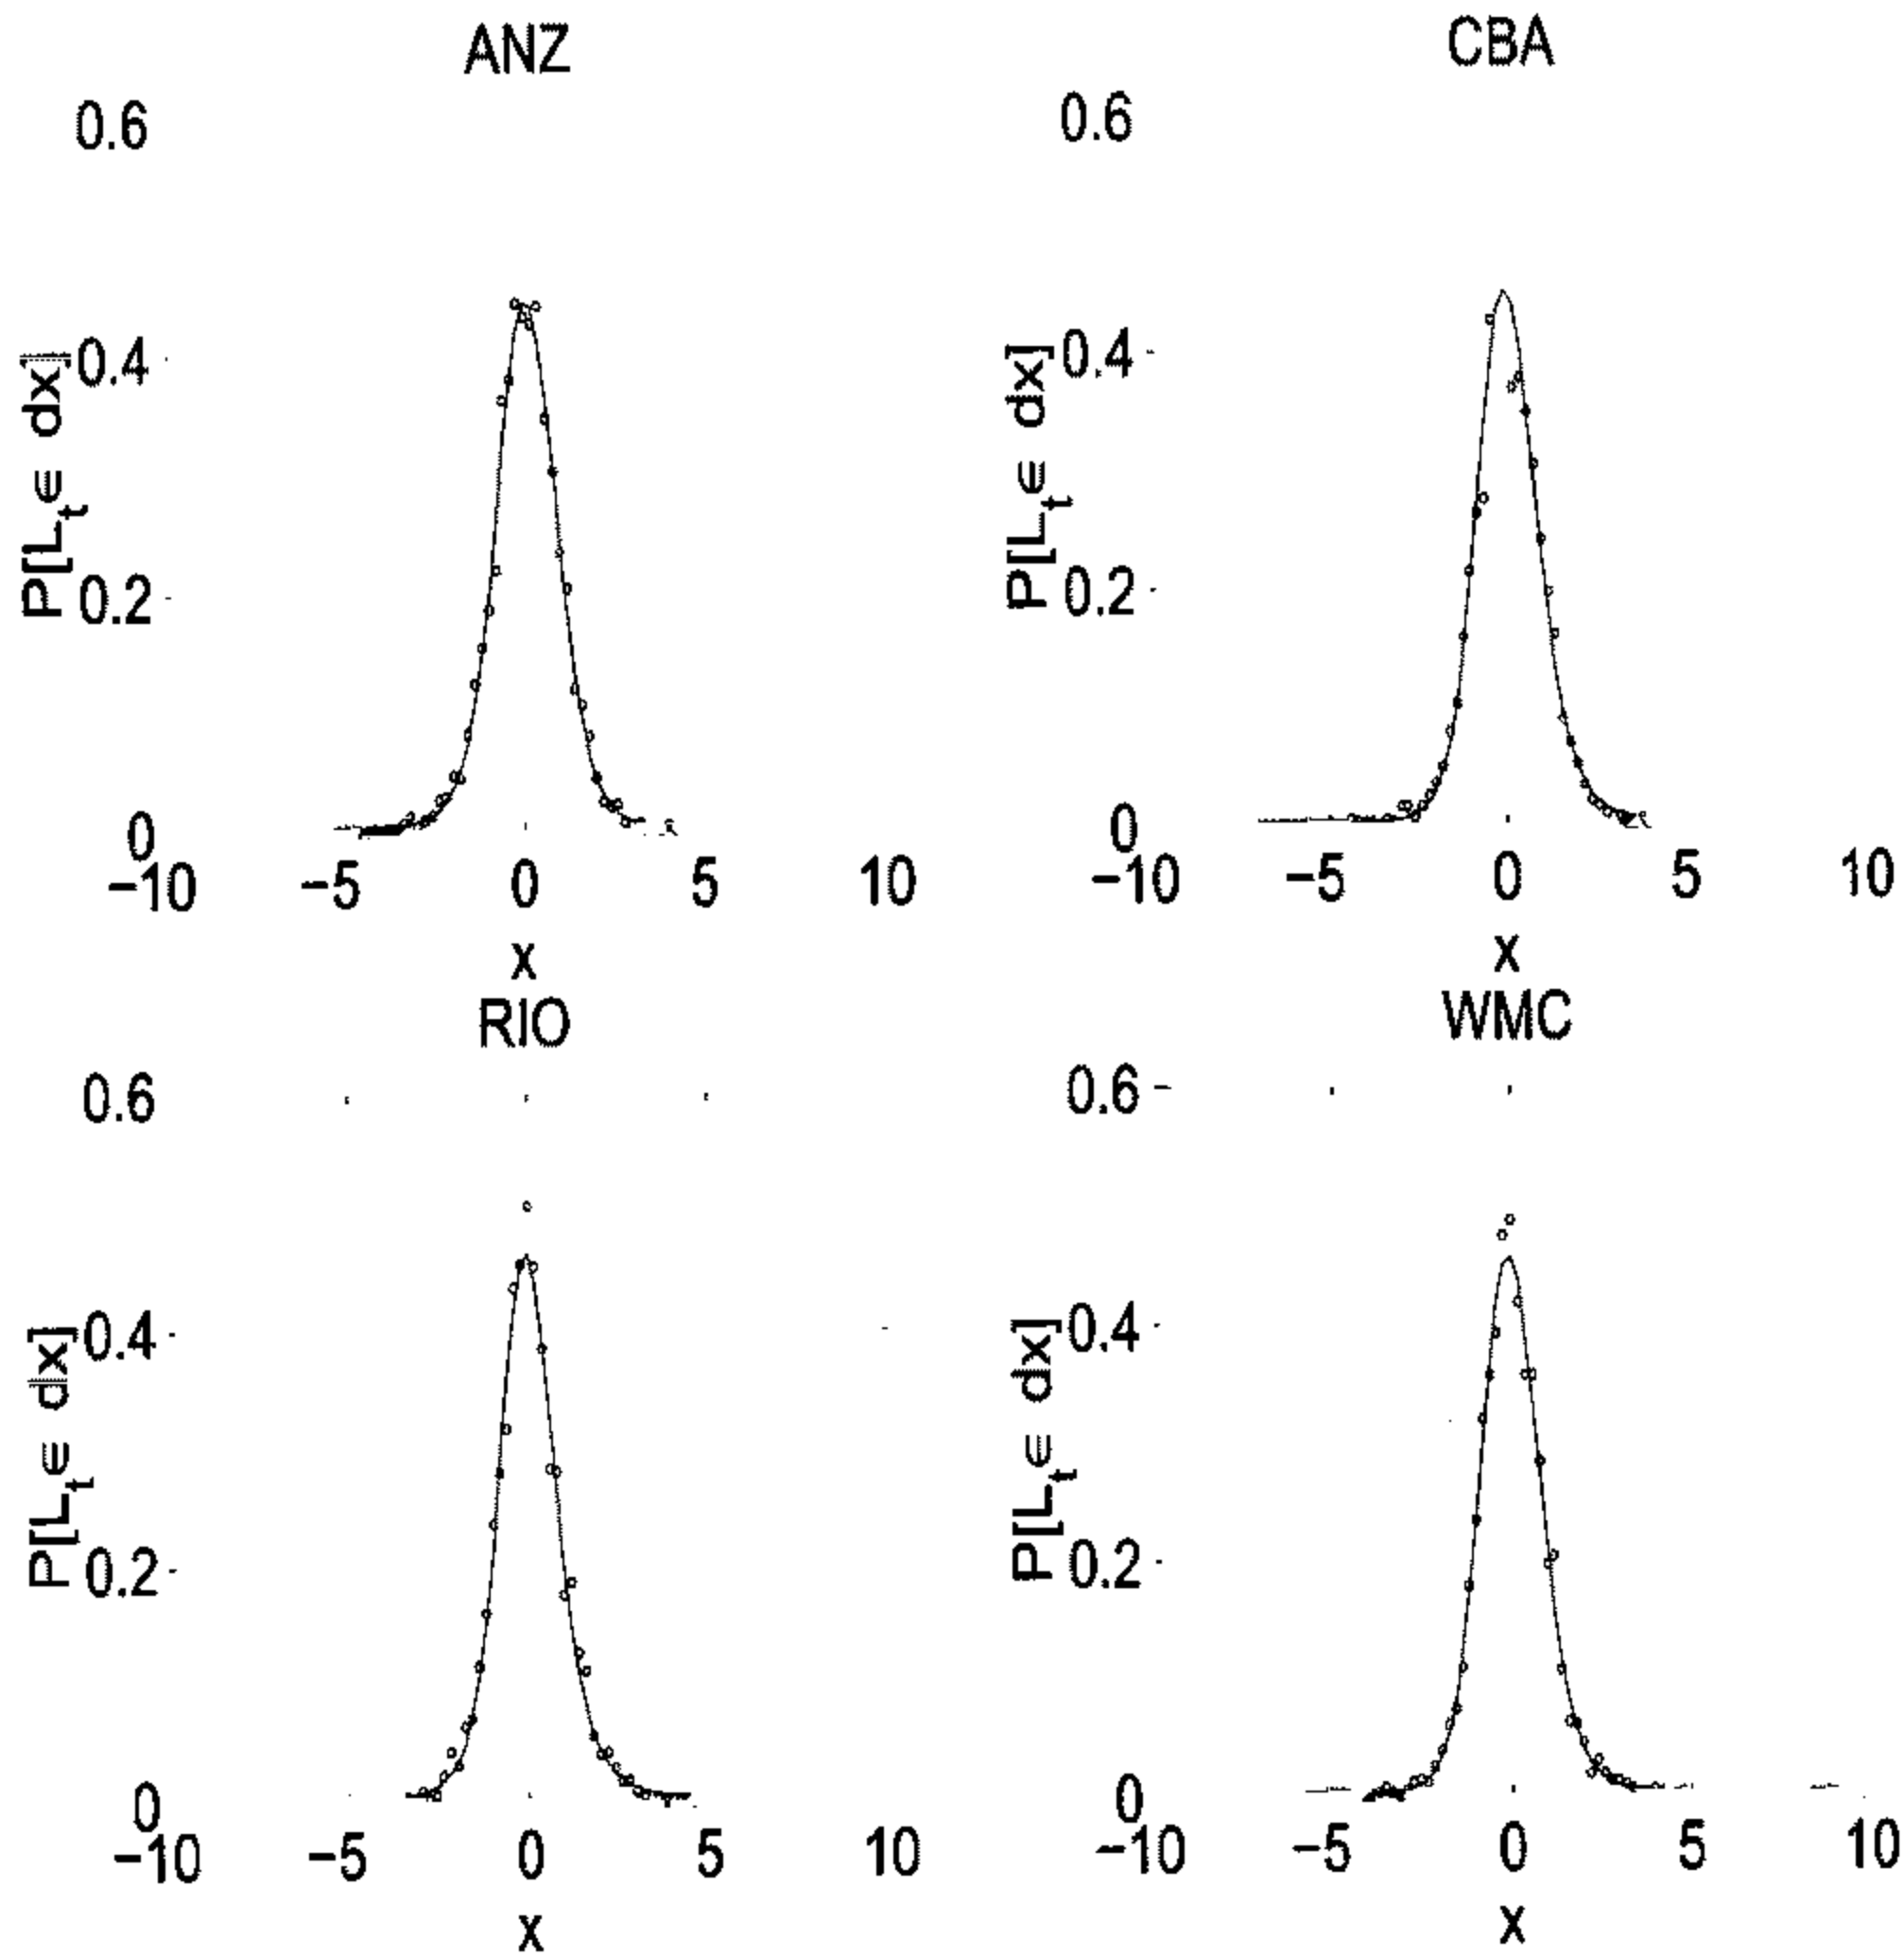


Figure 2.2: This figure shows the unconditional four parameter Lévy model density fitted to the daily log-return of four Australian stocks.

the fit concentrates on the central region of the density function rather than the distribution tails. Later, in Chapter 3 where we investigate the properties of high frequency data, we will explore further the tail behaviour of the Lévy-stable model.

The Lévy-stable model has both benefits and drawbacks as an asset model. Mandelbrot's hypothesis allows for a much more general model than that of Brownian motion. The model specifically allows for the description of heavy tailed distributions through the index parameter α . With this parameter the Lévy-stable model

is able to duplicate the extreme changes in prices that are so important for option pricing and risk management. The property of being closed under convolutions is often seen as being desirable from a theoretical point of view, as the distribution over any period of time will then be of the same class. However, this may or may not be the case with real data. It has been noted in more recent studies such as Gopikrishnan et.al. (1999) that the functional form of the distribution seems to change over time. The Lévy-stable model suffers from the drawback that log-returns are assumed to be independent identically distributed random variables, meaning that any possible dependence structure in the stock price is not addressed with this model. There have been a multitude of studies reporting on the correlation behaviour of log-return data and it is becoming widely agreed that independent identically distributed log-returns cannot be assumed a priori. Finally, despite what is said by supporters of this model, the infinite variance property of the α -stable Lévy process is a major drawback in the functionality of the model. Possessing no closed form for the distribution function is another detractor, leading to analytic and numerical complications.

2.3 Subordinated models

Subordinated asset models (or mixture models) were first proposed by Mandelbrot and Taylor (1967) and Clark (1973). This type of model offers a more physical description of the dynamic structure of asset returns than that obtained by simply fitting the distribution for the asset return. However, it arises that most of the models that fall into this class end up proposing unconditional return distributions simply to fit the data. The method of subordination goes back to the work of Bochner (1955) and was developed further in Feller (1966).

If W_t is a stochastic process and T_t is a non-negative, increasing stochastic process, we may form a new process $W(T_t)$. We say that the new process is subordinate to W_t by the driving process T_t . The method of subordination thus describes the behaviour of random events taking place on a random time scale. For a subordinated

model the asset price P_t is described by a stochastic equation of the form,

$$\log P_{t+\tau} - \log P_t = \mu\tau + \sigma (W(T_{t+\tau}) - W(T_t)) \quad (2.43)$$

where μ and σ are constant parameters and W_t is a stochastic process representing the random change in price. This method generalises the lognormal pricing model by introducing a random time scale T_t which specifies the intrinsic time scale on which the market evolves. If W_t is taken to be a Wiener process then the asset price viewed over time scale T_t will evolve according to a geometric Brownian motion, but when viewed in real time the asset price may behave quite differently. The evolution of the price process on this new time scale can be thought of in terms of information flow. As Clark (1973) writes

“The different evolution of price series on different days is due to the fact that information is available to traders at a varying rate. On days when no new information is available, trading is slow, and the price process evolves slowly. On days when new information violates old expectations, trading is brisk and the price process evolves much faster.”

Mixture models are thus able to explain the deviation from the normal distribution as seen in the data. The quantity of interest now becomes the random time scale T_t . Studies such as Mandelbrot and Taylor (1967), Clark (1973) and Geman and Ané (1996) have linked this stochastic clock with both the cumulative volume traded and the number of transactions. Many of the models in this class do not identify the process T_t with a measurable quantity and the choice of distribution for T_t appears to be specified somewhat arbitrarily.

In this section of Chapter 2 we review several of the many subordinated asset models. The models included in this section were chosen for their relative importance, new features, and popularity. We fit the models to Australian stock data and compare their relative performance. This list is by no means exhaustive and we must make note of other subordinated models: Barndorff-Nielsen (1995), Eberlien and Keller (1995) and Mittnik and Rachev (1993) although we do not specifically detail them here.

2.3.1 Properties of subordinated processes

Let W_t be a Wiener process representing the random effects in the evolution of a stock price. Let T_t be a random process, independent of W_t , with non-negative stationary independent increments, representing the speed of evolution. Then the log-return is given by

$$L_t = \mu\tau + \sigma [W(T_{t+\tau}) - W(T_t)]. \quad (2.44)$$

The process L_t has stationary independent increments given by the distribution function,

$$F_L(x) = \int_0^\infty \mathcal{N}\left(\frac{x - \mu\tau}{\sigma\sqrt{u}}\right) dF_T(u) \quad (2.45)$$

where $\mathcal{N}(x)$ is the normal distribution function and $F_T(t)$ is the distribution function of the process T_t . The probability density function for the log-return is given by

$$f_L(x) = \int_0^\infty \frac{1}{\sqrt{2\pi u\sigma}} e^{-\frac{(x-\mu\tau)^2}{2\sigma^2 u}} dF_T(u). \quad (2.46)$$

Recall that characteristic function of a random variable is defined as the Fourier transform of its distribution function,

$$\phi(\theta) = \int_0^\infty e^{i\theta x} dF(x) \quad (2.47)$$

then the characteristic function of a subordinated process can be written,

$$\phi_L(\theta) = \int_0^\infty e^{i\mu\tau\theta - \frac{1}{2}\sigma^2 u\theta^2} dF_T(u). \quad (2.48)$$

We can derive expressions for the k^{th} moment, m_k , of the return distribution using the well known formula,

$$m_k = \mathbb{E}_L[L_t^k] = (-i)^k \frac{\partial^k \phi_L}{\partial \theta^k} \Big|_{\theta=0} \quad (2.49)$$

where $E_L[\cdot]$ represents the expectation with respect to the distribution function of the variable L_t . Using Eq. (2.48) and Eq. (2.49), the first four moments are given by,

$$m_1 = \mu\tau \quad (2.50)$$

$$m_2 = \mu^2 \tau^2 + \sigma^2 \mu_T \quad (2.51)$$

$$m_3 = \mu^3 \tau^3 + 3\mu\sigma^2 \tau \mu_T \quad (2.52)$$

$$m_4 = \mu^4 \tau^4 + 6\mu^2 \sigma^2 \tau^2 \mu_T + 3\sigma^4 \sigma_T^3 \quad (2.53)$$

and the centralised moments,

$$\mu_L = \mathbb{E}[L_t] = \mu\tau \quad (2.54)$$

$$\sigma_L^2 = \mathbb{E}[(L_t - \mu_L)^2] = \sigma^2 \mu_T \quad (2.55)$$

$$\beta_L = \frac{\mathbb{E}[(L_t - \mu_L)^3]}{\sigma_L^3} = 0 \quad (2.56)$$

$$\kappa_L = \frac{\mathbb{E}[(L_t - \mu_L)^4]}{\sigma_L^4} = 3 \left(1 + \frac{\sigma_T^2}{\mu_T^2} \right) \quad (2.57)$$

where μ_T and σ_T^2 are the mean and variance of the random variable T_t . Eqs. (2.54-2.57) represent the mean, variance, skewness and kurtosis of the log-return process. The kurtosis parameter measures the tail thickness and peakedness of the distribution. The normal distribution has kurtosis $\kappa_L = 3$, while distributions with $k > 3$ are said to display leptokurtosis. When T_t is deterministic (with $\sigma = 0$) the increment of the log-asset price, L_t is normally distributed with mean $\mu_L = \mu\tau$, variance $\sigma_L^2 = \sigma^2\tau$ and kurtosis $\kappa_L = 3$. On the other hand, for any random time scale, the kurtosis will always display $\kappa_L > 3$ indicating fatter tails and higher peaks than that of the normal distribution. Thus, by introducing a stochastic time scale one naturally obtains a fat tailed distribution for the log-return. It can also be observed that the variance of a subordinated process will be finite only if the driving process has finite mean.

Using Eq. (2.47) and Eq. (2.48) we can express the characteristic function of the log-return in terms of the characteristic function of the driving process T_t ,

$$\phi_L(\theta) = e^{i\mu\tau\theta} \phi_T\left(i\frac{\sigma^2}{2}\theta^2\right). \quad (2.58)$$

This representation provides us with a way of determining the moments of L_t directly from those of T_t .

It is well known that any stochastic volatility model can be rewritten as a subordinated Brownian motion. The Dubins-Schwartz theorem (see Karatzas and Shreve

(1988)) allows for the representation of any Brownian semi-martingale as a time changed Brownian motion,

$$X_t = \int_0^t \sigma(t) dW_t \stackrel{d}{=} B(T_t) \quad \text{with} \quad T_t = \int_0^t \sigma^2(t) dt. \quad (2.59)$$

The above results will be used in the following sections to examine the properties of the subordinated models under study.

2.3.2 The Mandelbrot model

Mandelbrot and Taylor (1967) showed that the stable Paretian hypothesis could be interpreted as Brownian motion with a Lévy Stable function of time. For the subordinated return process to become the α -stable Lévy process of Mandelbrot (1963), the transition density of the driving process must be a scaled maximally skewed $\frac{\alpha}{2}$ -stable Lévy density. That is,

$$T_{t+\tau} - T_t \stackrel{d}{=} S_{\frac{\alpha}{2}}(c\tau^{2/\alpha}, 1, 0) \quad (2.60)$$

where the skewness parameter, $\beta = 1$. Note that because we require $0 < \alpha < 2$, the driving process will have an infinite mean since its has stability index $\frac{\alpha}{2} < 1$. Thus the resulting process will possess the infinite variance property.

The characteristic function of the time increment in Eq. (2.60) is given by

$$\phi_T(\theta) = \exp \left\{ -c^{\alpha/2} \tau |\theta|^{\alpha/2} \left(1 - i \tan \left(\frac{\pi\alpha}{4} \right) \text{sgn}(\theta) \right) \right\}. \quad (2.61)$$

Then using Eq. (2.58) we find that

$$\phi_L(\theta) = e^{i\mu\tau\theta} \phi_T\left(i\frac{\sigma^2}{2}\theta^2\right) \quad (2.62)$$

$$= \exp \left\{ i\mu\tau\theta - \frac{c^{\alpha/2}\tau |\sigma^2\theta^2|^{\alpha/2}}{2^{\alpha/2}} \left(1 + \tan \left(\frac{\pi\alpha}{4} \right) \text{sgn}(\theta^2)\text{sgn}(\sigma^2) \right) \right\}. \quad (2.63)$$

Since the term $\text{sgn}(\theta^2)\text{sgn}(\sigma^2) = 1$ unless $\theta = 0$ we can write the characteristic function as

$$\phi_L(\theta) = \exp \{ i\mu\tau\theta - \zeta^\alpha \tau |\theta|^\alpha \} \quad (2.64)$$

where

$$\zeta = \sqrt{\frac{c\sigma^2}{2}} \left(1 + \tan\left(\frac{\pi\alpha}{4}\right)\right)^{\frac{1}{\alpha}}. \quad (2.65)$$

This is the characteristic function for an α -stable Lévy motion with scale parameter $\zeta\tau^{1/\alpha}$. Hurst, Platen and Rachev (1997) and Rachev and Mittnik (2000) used the Laplace-Stieltjes transform to obtain this characteristic function and obtain a different representation of the parameter ζ . They mistakenly claimed that the value for ζ given in Mandelbrot and Taylor (1967) (and which is reproduced here) is erroneous.

Under the Mandelbrot and Taylor model the asset log-return is distributed as a stationary random variable,

$$L_t \sim S_\alpha(\zeta\tau^{1/\alpha}, 0, \mu\tau) \quad (2.66)$$

with infinite moments of index value greater than α . To characterise the directing process it is only necessary to specify the index of stability α . Further, Eq. (2.33) shows it is clear that as $\alpha \rightarrow 2$ the model reduces to the lognormal case.

Other models based on the stable Paretian hypothesis have been recently developed such as Mittnik and Rachev (1993). These models extend Mandelbrot's initial work to other more general stable processes such as the Laplace process.

2.3.3 The Clark lognormal model

Clark (1973) proposed a subordinated model where the price series evolves on a time scale determined by the lognormal distribution. Clark's model was designed to address the infinite variance feature of Mandelbrot's model. The lognormally driven geometric Brownian motion has finite moments of all orders. Clark argued that the central limit theorem holds only when the number of random variables occurring in each time step is almost constant and that speculative markets violate this property. In this case the limiting distribution will be subordinate to a normal. As a proxy for the driving process he suggested cumulative volume. Under this model the increments of time are independent random variables drawn from a lognormal distribution,

$$T_{t+\tau} - T_t \stackrel{d}{=} LN(e^{-\frac{1}{2}\chi^2}, \chi) \quad (2.67)$$

where LN denotes the lognormal distribution as given in Hastings and Peacock (1974). The first parameter is a scale parameter and the second is a shape parameter that reflects tail thickness. The increment $T_{t+\tau} - T_t$ has density function,

$$f_T(t) = \frac{1}{\chi\sqrt{2\pi}} t^{-1} e^{-\frac{(\log t + \frac{1}{2}\chi^2)^2}{2\chi^2}}; \quad t > 0 \quad (2.68)$$

such that $\mu_T = 1$ and $\sigma_T^2 = e^{\chi^2} - 1$. The resulting distribution function of the log-return admits the representation,

$$f_L(x) = \int_0^\infty \mathcal{N}\left(\frac{x - \mu\tau}{\sigma\sqrt{u}}\right) \frac{1}{\chi\sqrt{2\pi}} u^{-1} e^{-\frac{(\log u + \frac{1}{2}\chi^2)^2}{2\chi^2}} du \quad (2.69)$$

which has no closed form solution. By Eqs. (2.54-2.57) it is found that this model has mean $\mu_L = \mu\tau$, variance $\sigma_L^2 = \sigma^2\tau$ and kurtosis, $\kappa_L = 3(1 + e^{\chi^2} - 1)$. We can see here how the parameter χ controls the thickness of the tails through the kurtosis. As this parameter decreases, $\chi \rightarrow 0$, the time scale reduces to real time and the distribution becomes Gaussian.

Clark's model can be interpreted as a stochastic volatility model where the instantaneous volatility is lognormal. In fact, this model is quite similar to the Hull and White (1987) stochastic volatility model. One difference is that the Hull and White model is closed under convolutions whereas the Clark model is not. This means that we must identify a specific time scale on which to apply the model, as other time scales will be of differing functional forms.

Clark (1973) tested this model against the symmetric Lévy-stable model of Mandelbrot, reporting a better fit of the data. While Clark argued about the merits of finite variance, Mandelbrot argued that the costs of having a non-stable model outweighed the benefits of including finite moments. Both models have no simple analytic form for the transition density function. However, the Lévy-stable model benefits from admitting a relatively simple form for the characteristic function, while the lognormal model does not possess a similar property.

2.3.4 The Student- t model

The Lévy distribution possesses fat tails and infinite variance, while the Gaussian distribution has thin tails and finite moments of all orders. Lying between these two

extremes is the Student- t distribution. The Student- t distribution is specified in terms of the parameter $\nu > 0$, known as the degrees of freedom of the distribution. With this parameter it is possible to control the fatness of the distribution tails and the number of finite moments. If $\nu = 1$, we have the Lorentzian distribution with infinite mean. Conversely as $\nu \rightarrow \infty$ we approach the Gaussian distribution with finite moments of all orders. The Student- t distribution has a finite moment of order k for $k < \nu$. For $\nu \neq 1$ the distribution is not Lévy-stable and consequently distributions at different time horizons do not obey scaling relations.

Praetz (1969, 1972) found that under the random walk hypothesis, the sample variance, σ^2 , of real asset data appeared to vary significantly from year to year. Praetz (1972) proposed that the variance is not constant but can be specified by a random variable with distribution $g_\sigma(x)$. In the random walk Brownian motion model the log-return is given by a normal distribution. Introducing a random variance means that the log-return density function must be reinterpreted as a conditional distribution,

$$f_L(x) = \int_0^\infty f_L(x|\sigma^2)g_\sigma(\sigma^2)d\sigma^2 \quad (2.70)$$

where the probability density of the log-return for a given σ^2 , $f_L(x|\sigma^2)$, is Gaussian,

$$f_L(x|\sigma^2) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-\mu\tau)^2}{2\sigma^2}}. \quad (2.71)$$

Praetz (1971) suggested the variance be distributed as an inverted gamma random variable with density function,

$$g_\sigma(u) = \frac{(\nu/2)^{\nu/2}}{\Gamma(\nu/2)} u^{-\nu/2-1} e^{-\frac{\nu}{2u}}; \quad \nu \in \mathbb{R}^+. \quad (2.72)$$

Substituting this distribution into Eq. (2.70) and evaluating the integral we obtain,

$$f_L(x) = \int_0^\infty \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-\mu\tau)^2}{2\sigma^2}} \frac{(\nu/2)^{\nu/2}}{\Gamma(\nu/2)} \sigma^{-\nu-2} e^{-\frac{\nu}{2\sigma^2}} d\sigma^2 \quad (2.73)$$

$$= \frac{\Gamma((\nu+1)/2)}{\sigma\sqrt{\pi}\Gamma(\nu/2)} \left(1 + \frac{(x-\mu)^2}{\sigma^2\nu}\right)^{-(\nu+1)/2}. \quad (2.74)$$

The resulting distribution in Eq. (2.74) is the Student- t distribution with ν degrees of freedom and scale parameter σ . The Student- t model has mean $\mu_L = \mu\tau$, variance $\sigma_L^2 = \sigma^2 \frac{\nu}{\nu-2}\tau$ if $\nu > 2$ and kurtosis $\kappa_L = 3 \left(1 + \frac{2}{(\nu-4)\tau}\right)$ if $\nu > 4$.

Praetz (1972) studied weekly Australian equity data and found that the Student $-t$ model out-performed the Lévy-stable model. Blattberg and Gonendes (1974) compared both the Student and the Lévy-stable models to daily US asset data and found that the Student model described the data much better than the Lévy-stable model. However, Praetz's analysis was criticised for normalising the data by its standard deviation, while at the same time fitting an infinite variance stable distribution to the data. Blattberg and Gonendes (1977) and Praetz (1977) argued the significance of using the sample standard deviation as a measure for scale of the non-Gaussian stable distributions. Though the Student- t distribution is infinitely divisible, and hence consistent with a process of stationary independent increments, its analytical structure is difficult to work with.

2.3.5 The variance gamma model

The variance gamma (VG) model is a popular model that still attracts a great deal of interest from researchers and practitioners. Madan and Seneta (1990) developed a new stochastic process for asset returns known as the VG process. In their opinion it is more useful to propose a stochastic process rather than a distribution as this allows for applications such as option pricing through the identification of an explicit change of measure.

This model sought to describe the stylised features of equity data while at the same time address the deficiencies that affected other models. They set out to include:

- fat tailed return distributions that approach normality over large scales,
- finite moments for at least the lower orders of the log-return,
- consistency with an underlying continuous-time process with independent identically distributed increments which belong to the same distribution family, irrespective of interval length,
- extension to multivariate processes in order to preserve the capital asset pricing model (CAPM).

In the variance gamma model the log-return, L_t is normally distributed with mean $\mu_L = \mu$ and random variance $\sigma_L^2 = \sigma^2 V$. The distribution of V is taken to be the Gamma distribution with density function,

$$g_V(u) = \frac{c^\gamma u^\gamma e^{-cu\gamma}}{\Gamma(\gamma)} \quad c, \gamma > 0 \quad (2.75)$$

thus the log-return density function can be written

$$f_L(x) = \int_0^\infty \frac{1}{\sqrt{2\pi u\sigma}} e^{-\frac{(x-\mu)^2}{2\sigma^2 u}} \frac{c^\gamma u^\gamma e^{-cu\gamma}}{\Gamma(\gamma)} du. \quad (2.76)$$

Again, like the Student- t model there exists no closed form for this integral. However, it is possible to obtain an asymptotic representation for large x ,

$$f_L(x) = \frac{\sqrt{2/v} (x\sqrt{2/v}/\sigma)^{(2/v-1)/2}}{\sigma 2^{(2/v-1)/2} \Gamma(1/v) \sqrt{\pi}} K_{(2/v-1)/2}(x\sqrt{2/v}/\sigma); \quad x \rightarrow \infty \quad (2.77)$$

where $K_n(x)$ is the modified Bessel function of the second kind of order w . Here $m = \gamma/c$ is the mean of the gamma density and $v = \gamma/c^2$ is its variance. The characteristic function for the log-return has the simple representation,

$$\phi_L(\theta) = e^{i\mu\tau\theta} \left(1 + \frac{\sigma^2\theta^2 v}{2m} \right)^{-\frac{m^2}{v}}. \quad (2.78)$$

Viewing the variable V as a random time change and assuming the expected random time change, $\mathbb{E}[T_t]$, is one unit of time we can set $m = 1$. This yields the characteristic function of the return distribution over the unit interval,

$$\phi_L(\theta) = e^{i\mu\tau\nu} \left(1 + \frac{\sigma^2\theta^2 v}{2} \right)^{-\frac{1}{v}}. \quad (2.79)$$

Using this characteristic function to calculate the moments, we find that the process has mean $\mu_L = \mu\tau$, variance $\sigma_L^2 = \sigma^2\tau$ and kurtosis $\kappa_L = 3(1+v)$. Thus v may be regarded as a measure of tail thickness.

The VG model is closed under convolutions and possesses finite moments of all orders. According to Madan and Seneta (1990) the rate of tail decay is exponential. The VG process can be expressed as the difference of two independent identically distributed gamma increments allowing us to write,

$$L_t = U_t - D_t \quad (2.80)$$

where U_t and D_t are processes of independent gamma increments with means $\mu_{U,D} = (\sigma/\sqrt{2v})h$ and variance $\sigma_{U,D} = \sigma^2\Delta t/2$ for interval length Δt . In this representation the normalised log-return, $\bar{L}_t = L_t/\sigma\sqrt{\Delta t}$, has characteristic function,

$$\phi_L = \left(1 + \frac{vu^2}{2\Delta t}\right)^{-\frac{\Delta t}{v}}. \quad (2.81)$$

Examining the large Δt behaviour of the above equation, it is clear that as $\Delta t \rightarrow \infty$ the characteristic function, and hence the distribution, of L_t approaches normality.

All of the processes under this model are Lévy processes (Madan and Seneta (1990)). The variance gamma process can be shown to be a pure jump process which can be approximated by a compound Poisson process with high jump frequency. Madan and Seneta (1987a) reported a good fit of the model to data.

The work of Madan and Seneta (1990) has been added to by Carr, Geman, Madan and Yor (2002) proposing the so called CGMY model. This model extends the variance gamma model allowing for more general Lévy processes with finite or infinite variation in order to add additional mathematically appealing features.

2.3.6 Comparison of models

In this section we fit the previously described subordinated models, including geometric Brownian motion, to Australian Stock Exchange daily equity data. Our purpose here is to offer a comparison of the different models and discover which one best describes ASX daily log-return data. In order to estimate the parameters for these models, we follow the method of Hurst, Platen and Rachev (1997) and employ unconstrained maximum likelihood estimation (MLE). Using this estimation procedure, the number of degrees of freedom of a model will be given by the total number parameters of the model. Since the subordinated models studied here are three parameter models we expect them to perform better than the two parameter geometric Brownian motion. We rate each model's performance based on the value of the log-likelihood function. The best model will be identified as displaying the largest log-likelihood value.

The data we use consist of the last traded price for a selection of equities traded on the ASX during the period 15/7/93 to 28/6/02. Fig. 2.3 and Fig. 2.4 display the

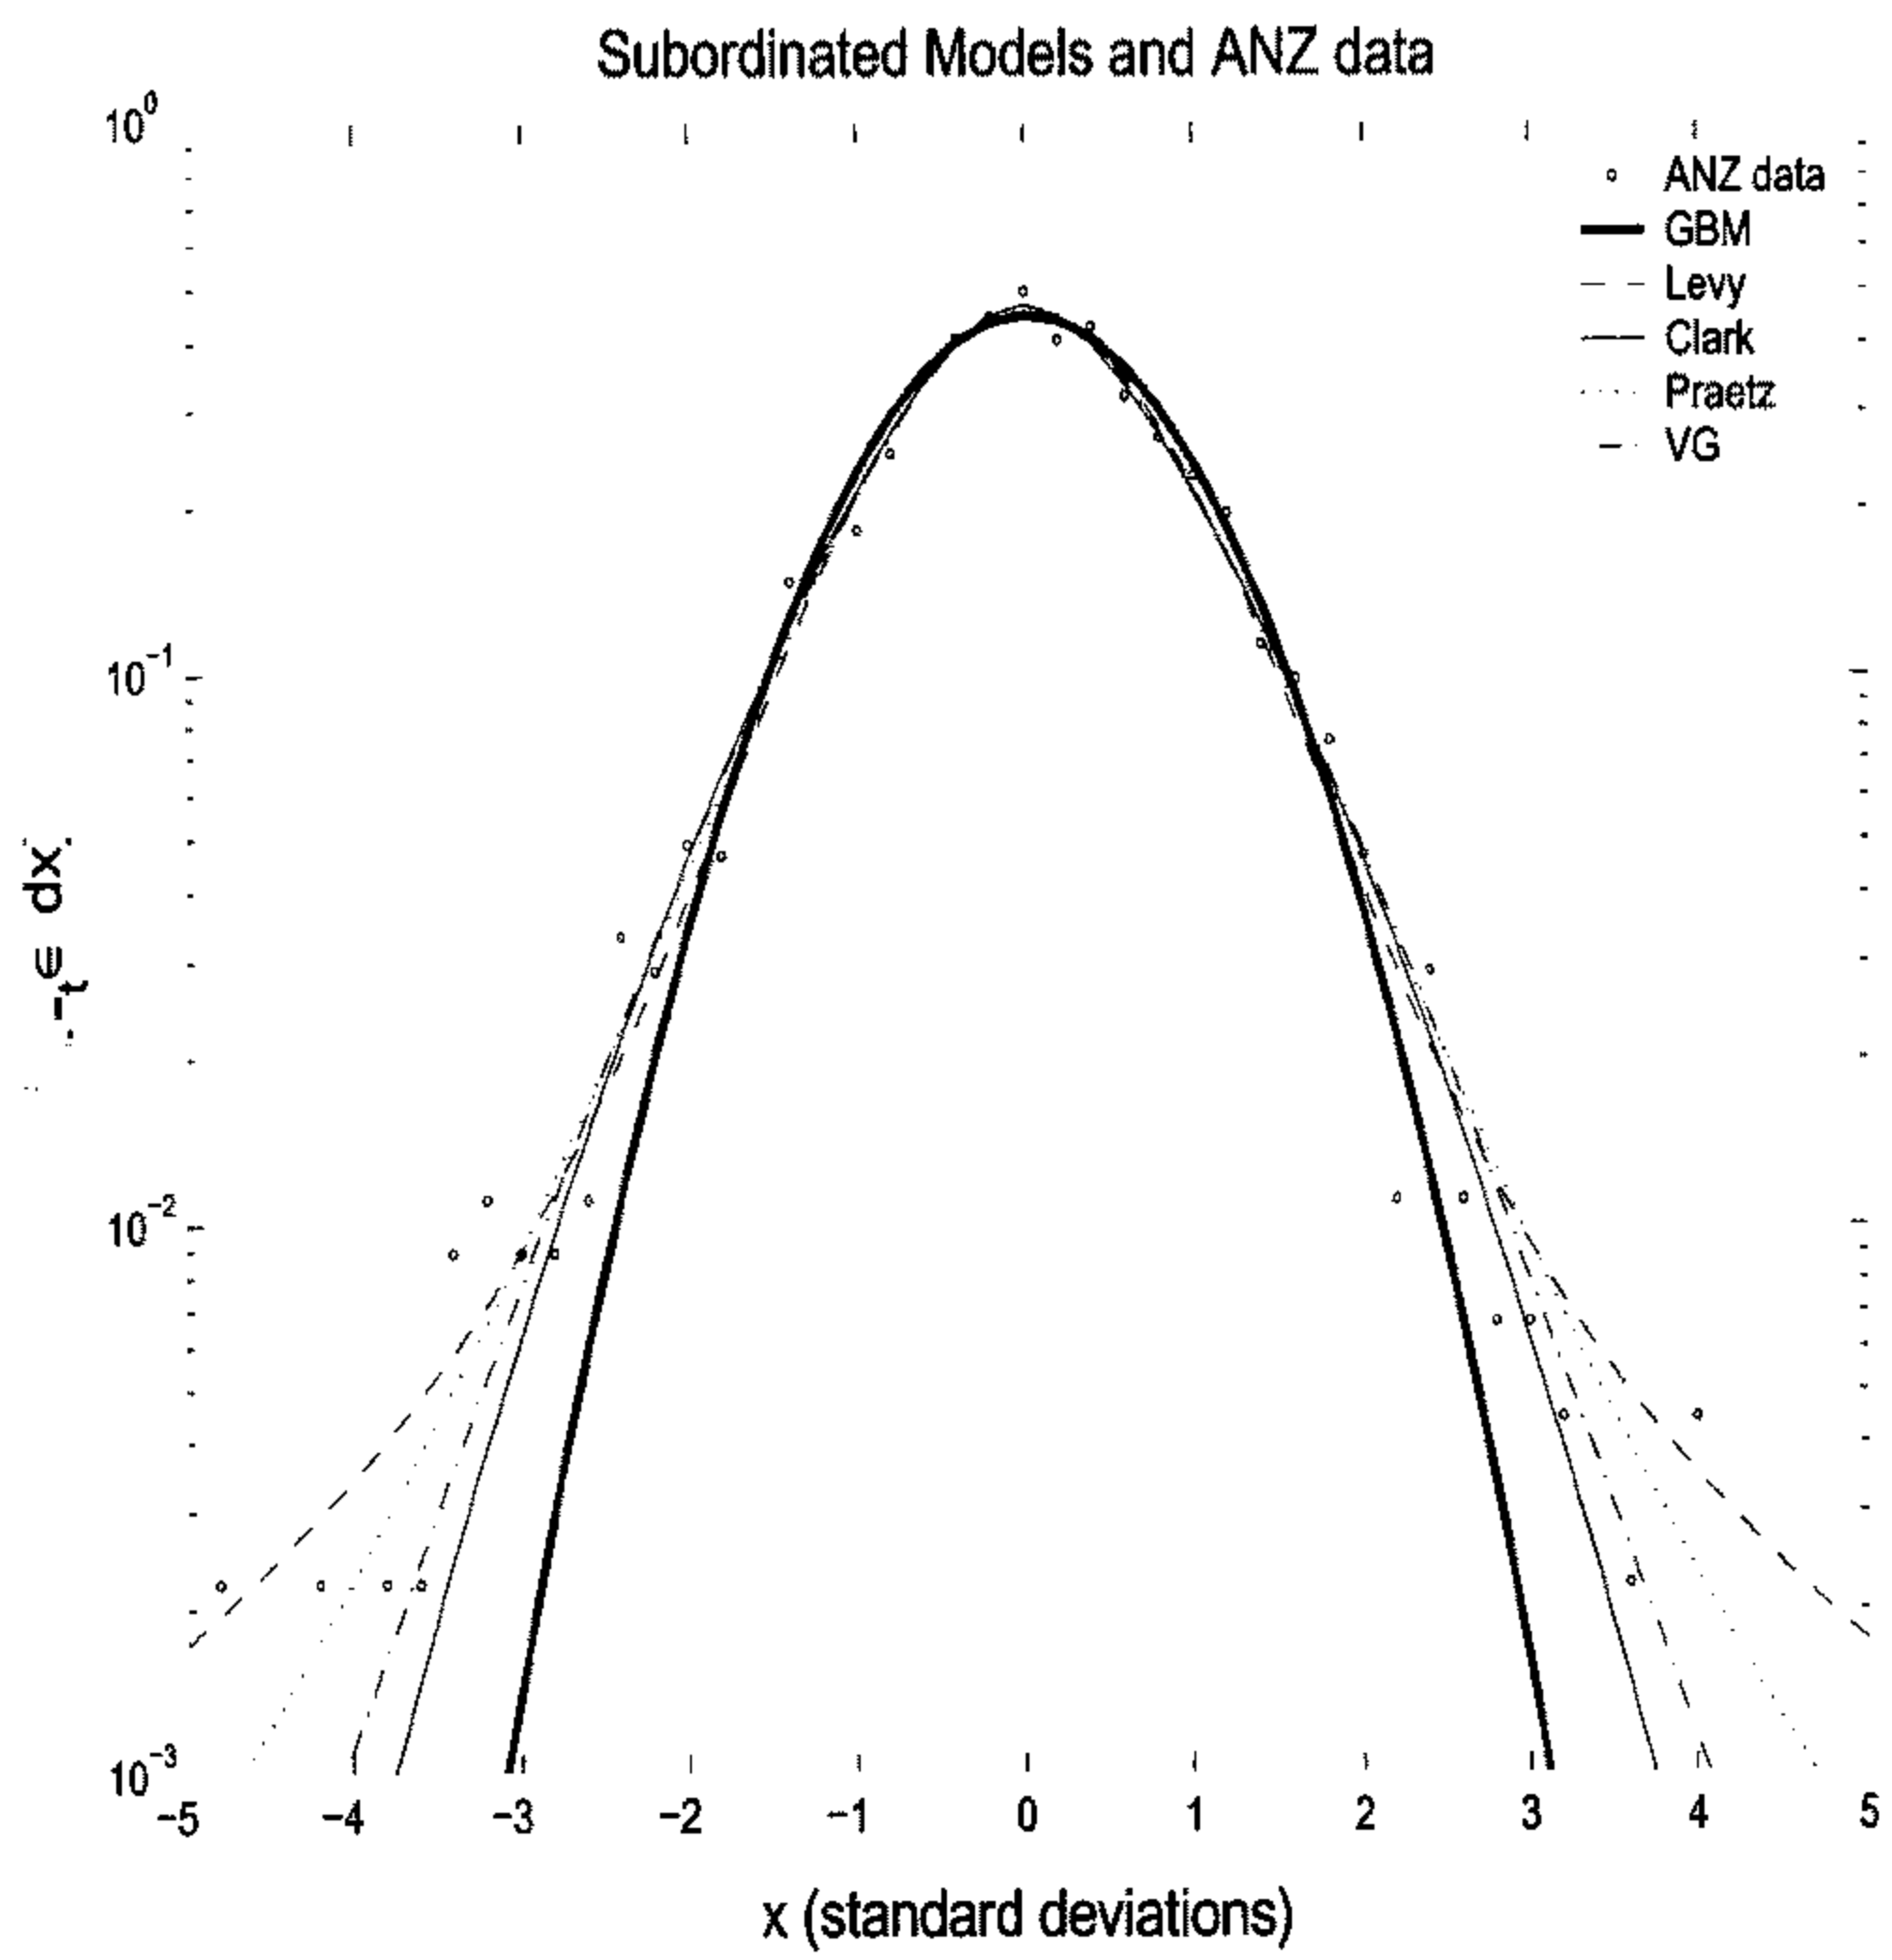


Figure 2.3: The estimated density function estimates for ANZ daily price data. The best model for this data is found to be the Student- t (Praetz) model with $\nu = 5.3$ degrees of freedom.




Figure 2.4: The estimated density functions for CBA daily price data. The best model for this data is found to be the Student- t (Praetz) model with $\nu = 6.9$ degrees of freedom.

density functions for the models estimated from ANZ and CBA daily return data. The models all show a marked improvement over that of geometric Brownian motion. Examination of the tails of the density functions show how the descriptive properties of the models differ. Notice that because we are using low frequency data, there are too few samples to establish a clear picture of the tail behaviour of empirical return data. This result points out a major problem that has plagued past studies in this field. The distribution tails represent the infrequent, large price changes and such inexact results mean that it is difficult for conclusions to be drawn as to what the correct model for the data is. Further, it is also this lack of tail data that has led to much of the controversy over the finiteness of moments. Tables A.2-A.6 contain the parameter values obtained by fitting the models to ten Australian Stocks. Table A.7 displays the models' log-likelihood functions for the parameters in Tables A.2-A.6. We see from this table that the Student- t model consistently displays the largest value for the log-likelihood function followed closely by the variance gamma model. The Student- t model estimates the number of degrees of freedom as $\nu \approx 4$ for most of the stocks listed in Table A.6, consistent with a finite fourth moment for the log-return. The Lévy-stable model returns $\alpha \approx 1.7$ indicating an infinite variance. The models of Clark and Mandelbrot rank third and fourth respectively.

2.4 Advanced subordinated models

All the models examined so far have focused solely on fitting the unconditional distribution of the log-return. These models all contained the requirement of independent identically distributed increments. Again, these same models ignored the fact that although the log-return itself seems to be independent, its higher order correlations indicate a non-negligible correlation structure.

Several studies such as Granger and Ding (1996) have investigated the presence of correlations and dependency in stock prices. Fig. 2.5 displays the autocorrelation function for the log-return, as well as the autocorrelation for the absolute and squared log-return. While the log-return itself appears to be uncorrelated, the absolute and squared log-return appears to display a distinct correlation.

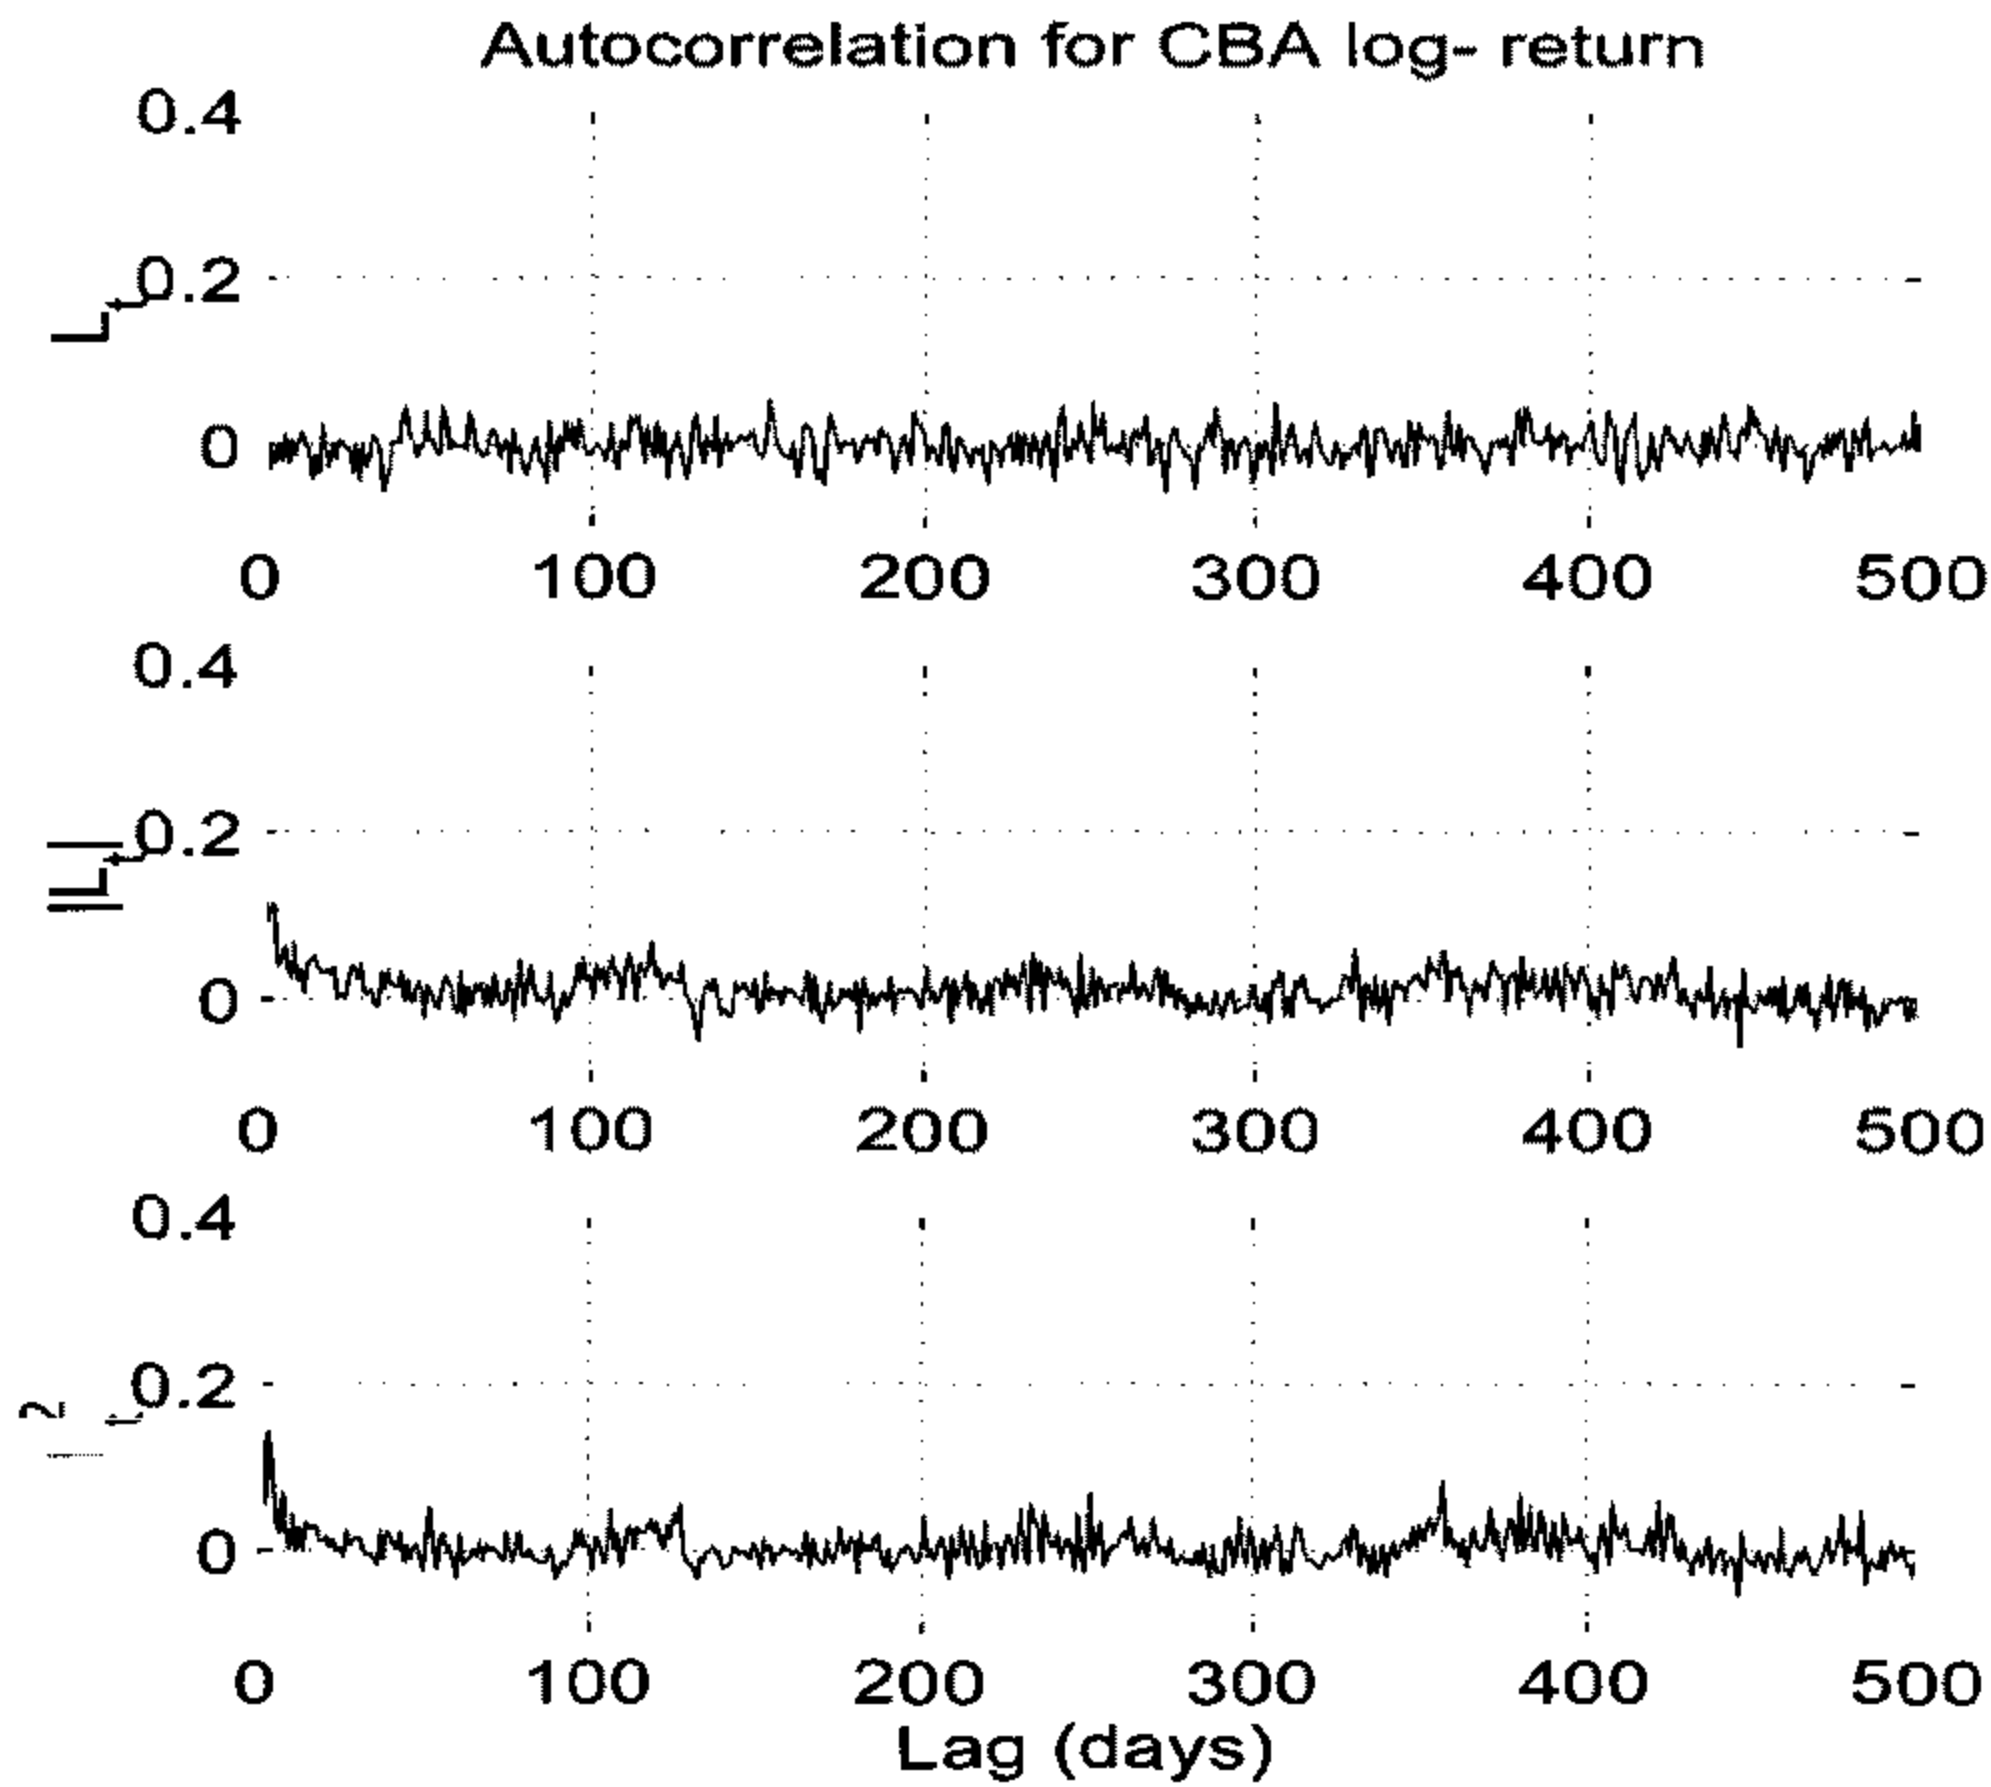


Figure 2.5: This figure shows the sample autocorrelation for the log-return, L_t , absolute log-return, $|L_t|$, and squared log-return, L_t^2 , for a typical ASX equity. The top plot seems to indicate uncorrelated random variables while the lower two plots display a slowly decreasing correlation structure.

This section details the more recent developments in the technique of subordination, namely, allowing the driving process to possess a dependence structure. With this more general form of compounding, the next three models we look at are able to address the temporal dependence structure of asset returns.

2.4.1 The Poisson process

The work of Edelman and Gillespie (1997) extends the method of stochastic subordination, assuming a suitable process for the driving process without having to identify it with a physically observed variable. The authors suggest that whatever the underlying measure of trading activity T_t , it should be represented as an inhomogenous Poisson process,

$$f_T(t) = \frac{\Lambda(t)^k}{k!} e^{-\Lambda(t)} \quad (2.82)$$

where

$$\Lambda(t) = \int_0^t \lambda(s) ds \quad (2.83)$$

and $\lambda(t)$ denotes the intensity of the process.

The Poisson process is a well known stochastic process, commonly used to model systems of randomly occurring events such as queuing and traffic flow. Edelman and Gillespie (1997) justify this choice by arguing that over a small time interval $[t, t + \epsilon]$ the occurrence of a trading event may be regarded as a rare event and hence modelled by a Poisson random variable with intensity $\lambda(t)$. From an applied mathematics perspective, Edelman and Gillespie's proposal of the Poisson process seems to be a more sensible choice for modelling the activity time than the assumption of a lognormal or $\alpha/2$ -stable random walk.

The intensity $\lambda(t)$ is taken to be a two state Markov process (λ_1, λ_2) , representing 'quiet' and 'busy' states of the market. Discretising time, the process $\lambda(t) = \lambda(0), \lambda(\tau), \lambda(2\tau) \dots$ can be thought of as a Markov Chain with transition matrix,

$$\mathbf{P} = \begin{pmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{pmatrix} \quad (2.84)$$

where

$$p_{ij} = \mathbf{P}[\lambda(t + \tau) = \lambda_j | \lambda(t) = \lambda_i] \quad (2.85)$$

are the probabilities of moving from state i to state j . The long run probabilities of being in one state or the other are found to be

$$p_1 = \frac{1 - p_{22}}{2 - p_{11} - p_{22}} \quad (2.86)$$

$$p_2 = 1 - p_1. \quad (2.87)$$

Under this model the density function of log-return is given by

$$f_L(x) = \sum_{i=1}^2 p_i \sum_{j=0}^{\infty} \frac{(\tau \lambda_i)^j}{j!} e^{-\tau \lambda_i} \frac{1}{\sqrt{2\pi j} \sigma} e^{-\frac{(x-\mu\tau)^2}{2\sigma^2 j}}. \quad (2.88)$$

When $j = 0$ the normal distribution approaches the Dirac delta function yielding,

$$f_L(x) = \sum_{i=1}^2 p_i e^{-\tau \lambda_i} \left(\delta(x - \mu\tau) + \sum_{j=0}^{\infty} \frac{(\tau \lambda_i)^j}{j!} \frac{1}{\sqrt{2\pi j} \sigma} e^{-\frac{(x-\mu\tau)^2}{2\sigma^2 j}} \right). \quad (2.89)$$

This density function can be approximated by the four term Edgeworth expansion,

$$f_L(x) = \sum_{i=1}^2 p_i e^{-\tau \lambda_i} \frac{e^{-\frac{(x-\mu)^2}{2\sigma^2 \lambda_i \tau}}}{\sqrt{2\pi \lambda_i \tau} \sigma} \left(1 + \frac{(r - \mu\tau)^4 - 6\sigma^2 \lambda_i \tau (r - \mu\tau)^2 + 3\sigma^4 \lambda_i^2 \tau^2}{8\sigma^4 \lambda_i^3 \tau^3} \right). \quad (2.90)$$

The resulting log-return process has mean, variance and kurtosis,

$$\mu_L = \mu\tau \quad (2.91)$$

$$\sigma_L^2 = \sigma^2 \tau (p_1 \lambda_1 + (1 - p_1) \lambda_2) \quad (2.92)$$

$$\kappa_L = 3(p_1 \lambda_1 (1 + \tau \lambda_1) + (1 - p_1) \lambda_2 (1 + \tau \lambda_2)). \quad (2.93)$$

Also, due to the Markov nature of the intensity $\lambda(t)$, the one time lag autocorrelation for the squared log-return is given by,

$$\frac{\mathbb{E}[R_t^2 R_{2t}^2]}{\mathbb{E}[R_t^2 - \mathbb{E}[R_t^2]]^2} = \frac{\sigma^4 t^2 (\lambda_2 - \lambda_1) p_1 (p_{11} - p_1)}{3\sigma^4 t (p_1 \lambda_1 (t\lambda_1 - 1) + (1 - p_1) \lambda_2 (t\lambda_2 - 1)) + 4\mu^2 \sigma^2 t^3 \bar{\lambda} - \sigma^4 t^2 \bar{\lambda}^2} \quad (2.94)$$

where $\bar{\lambda} = (p_1 \lambda_1 + (1 - p_1) \lambda_2)$.

This model is practically appealing as it can model the non-normality of returns, autocorrelation in squared returns and it is closed under convolutions. Edelman and Gillespie (1997) fitting the model to real data, found that it mimicked the deviations from normality as well as the observed autocorrelation structure. In order to improve the fit to real data they suggest that the Wiener process be replaced by a negatively skewed stable process.

2.4.2 FATGBM

Heyde (1999) provides a minimal description model that addresses the issue of heavy tails and the correlation structure of asset returns. Under this model, known as fractal activity time geometric Brownian motion (FATGBM), the log-return itself is uncorrelated while the absolute log-return and squared log-return both display a distinct correlation structure.

The model retains the GBM assumption that the log-return should have stationary increments, adding the requirements of heavy tails and strong dependence in the absolute and squared log-return. Heyde (1999) assumes that the activity time, T_t is a positive increasing random process independent of the Wiener process W_t . The activity time is assumed to have finite mean while its increments display a strong correlation structure.

The FATGBM model assumes that the price follows a stochastic process of the form,

$$P_t = P_0 e^{\mu t + \sigma W(T_t)}. \quad (2.95)$$

Using the self-similarity and stationary properties of the Wiener process, Eq. (2.9) and Eq. (2.14), and the log-return representation of Eq. (2.44) we have

$$L_t = \log P_{t+\tau} - \log P_t \quad (2.96)$$

$$= \mu\tau + \sigma(W(T_{t+\tau}) - W(T_t)) \quad (2.97)$$

$$\stackrel{d}{=} \mu\tau + \sigma\sqrt{T_{t+\tau} - T_t} W(1) \quad (2.98)$$

where $W(1) \sim N(0, 1)$ is a normally distributed random variable with zero mean and unit variance. Investigating the correlation structure of this process we find that,

$$\text{cov}(L_t, L_{t+k}) = 0 \quad (2.99)$$

$$\text{cov}(|L_t|, |L_{t+k}|) = \frac{2}{\pi} \text{cov}(\sqrt{T_t}, \sqrt{T_{t+k}}) \quad (2.100)$$

$$\text{cov}(L_t^2, L_{t+k}^2) = \sigma^4 \text{cov}(T_t, T_{t+k}) \quad (2.101)$$

where we have written $\mathcal{T}_t = T_{t+\tau} - T_t$ to represent the increments of the activity time. The dependence structure of the absolute and squared log-return can be deduced from the behaviour of $\sqrt{T_t}$ and \mathcal{T}_t respectively.

It is well known that Brownian motion is homoscedastic *i.e.* constant variance, since its conditional variance is

$$\mathbb{V}[L_t | \mathcal{F}_{t-\tau}] = \sigma^2 \tau \quad (2.102)$$

where $\mathcal{F}_{t-\tau}$ is the filtration consisting of the entire history of L_t up until $t - \tau$. In contrast, FATGBM is a heteroskedastic process *i.e.* changing variance, which varies through the conditional mean of the activity time increments,

$$\mathbb{V}[L_t | \mathcal{F}_{t-\tau}] = \sigma^2 \mathbb{E}[T_t | \mathcal{F}_{t-\tau}]. \quad (2.103)$$

Also the kurtosis of the model exceeds that of GBM,

$$\kappa_L = 3(1 + \mathbb{V}[T_t]) > 3. \quad (2.104)$$

Thus it can be seen that the dependence properties and distribution tail properties of the log-return are all derived from the properties of the activity time process and its increments.

Heyde (1999) suggests that the process $T_t - t$ is a self-similar process with scaling index H such that,

$$T_{ct} - ct \stackrel{d}{=} c^H (T_t - t) \quad (2.105)$$

where $\frac{1}{2} < H < 1$. It is this property which gives the model its name. Applying Itô's formula together with Eq. (2.95) allows one to obtain the expression,

$$d \log P_t - \frac{dP_t}{P_t} = \frac{1}{2} \sigma^2 dT_t. \quad (2.106)$$

Through this equation we are able to construct empirically the activity time process. Discretising and solving this equation for T_t we obtain a picture of the activity time. Fig. 2.6 displays both real time and activity time for several stocks. These plots show the extent to which the activity time deviates from that demanded by GBM. Heyde (1999) and Heyde and Liu (2001) explore the fractal nature of the activity time with an examination of empirical data. Writing

$$S_q(\delta) = \sum_{i=1}^N \frac{|T_{i\delta} - T_{(i-1)\delta} - \delta|^q}{N} \quad (2.107)$$

then if $T_t - t$ is self-similar,

$$S_q(\delta) \stackrel{d}{=} \delta^{qH} \sum_{i=1}^N \frac{|T_i - T_{(i-1)} - 1|^q}{N} \quad (2.108)$$

$$= \delta^{qH} S_q(1) \quad (2.109)$$

$$\Rightarrow \log \mathbb{E}[S_q(\delta)] = qH \log \delta + \log \mathbb{E}[S_q(1)]. \quad (2.110)$$

Using Eq. (2.110) we plot $\frac{\log \mathbb{E}[S_q(\delta)] - \log \mathbb{E}[S_q(1)]}{qH \log \delta}$ against q so that the slope will approximate² H . In Fig. 2.7 we see the results obtained from Eq. (2.110). It is clear that while these plots display a linear trend for $q < 1$, they show definite non-linear behaviour for larger values of q and hence cast some doubt on Heyde's assumption of the self-similarity of $T_t - t$.

FATGBM has been applied to the pricing of options and the calculation of value at risk (VaR). Heyde and Gay (2000) build on the excellent fit afforded by Praetz's Student- t model by suggesting that if T_1 is given by an inverted gamma variable,

$$f_{T_1}(t) = \frac{\left(\frac{\nu-2}{2}\right)^{\frac{\nu}{2}}}{\Gamma\left(\frac{\nu}{2}\right)} t^{-\frac{\nu}{2}-1} e^{-(\nu-2)/2t} \quad (2.111)$$

the activity process can be constructed via the self-similarity assumption of Eq. (2.105) with $c = t$ yielding,

$$T_t \stackrel{d}{=} t + t^H(T_1 - 1); \quad T_0 = 0. \quad (2.112)$$

However, if $t < 1$ and $0 < T_1 < 1 - t^H$, the activity time can be negative, $t - t^H < T_t < 0$. Hence restricting T_t to positive values, the distribution of T_t has an atom of probability at zero corresponding to the probability that T_1 lies in the domain $0 < T_1 < 1 - t^H$.

Heyde's FATGBM model possesses many desirable features for an asset model. It offers a minimal description model focusing on a simple form based on a minimum number of parameters designed to match stylised features of the data. Simulation is relatively straight forward and hence Monte Carlo approaches to option pricing

²This statistic is slightly different from the one that appears in Heyde's paper. A Sydney University School of Mathematics and Statistics report by Max Skipper finds Heyde's statistic to be missing a non-zero term.




Figure 2.6: The dotted line represents the real time scale under geometric Brownian motion while the solid line represents the actual activity time extracted via Eq. (2.106)

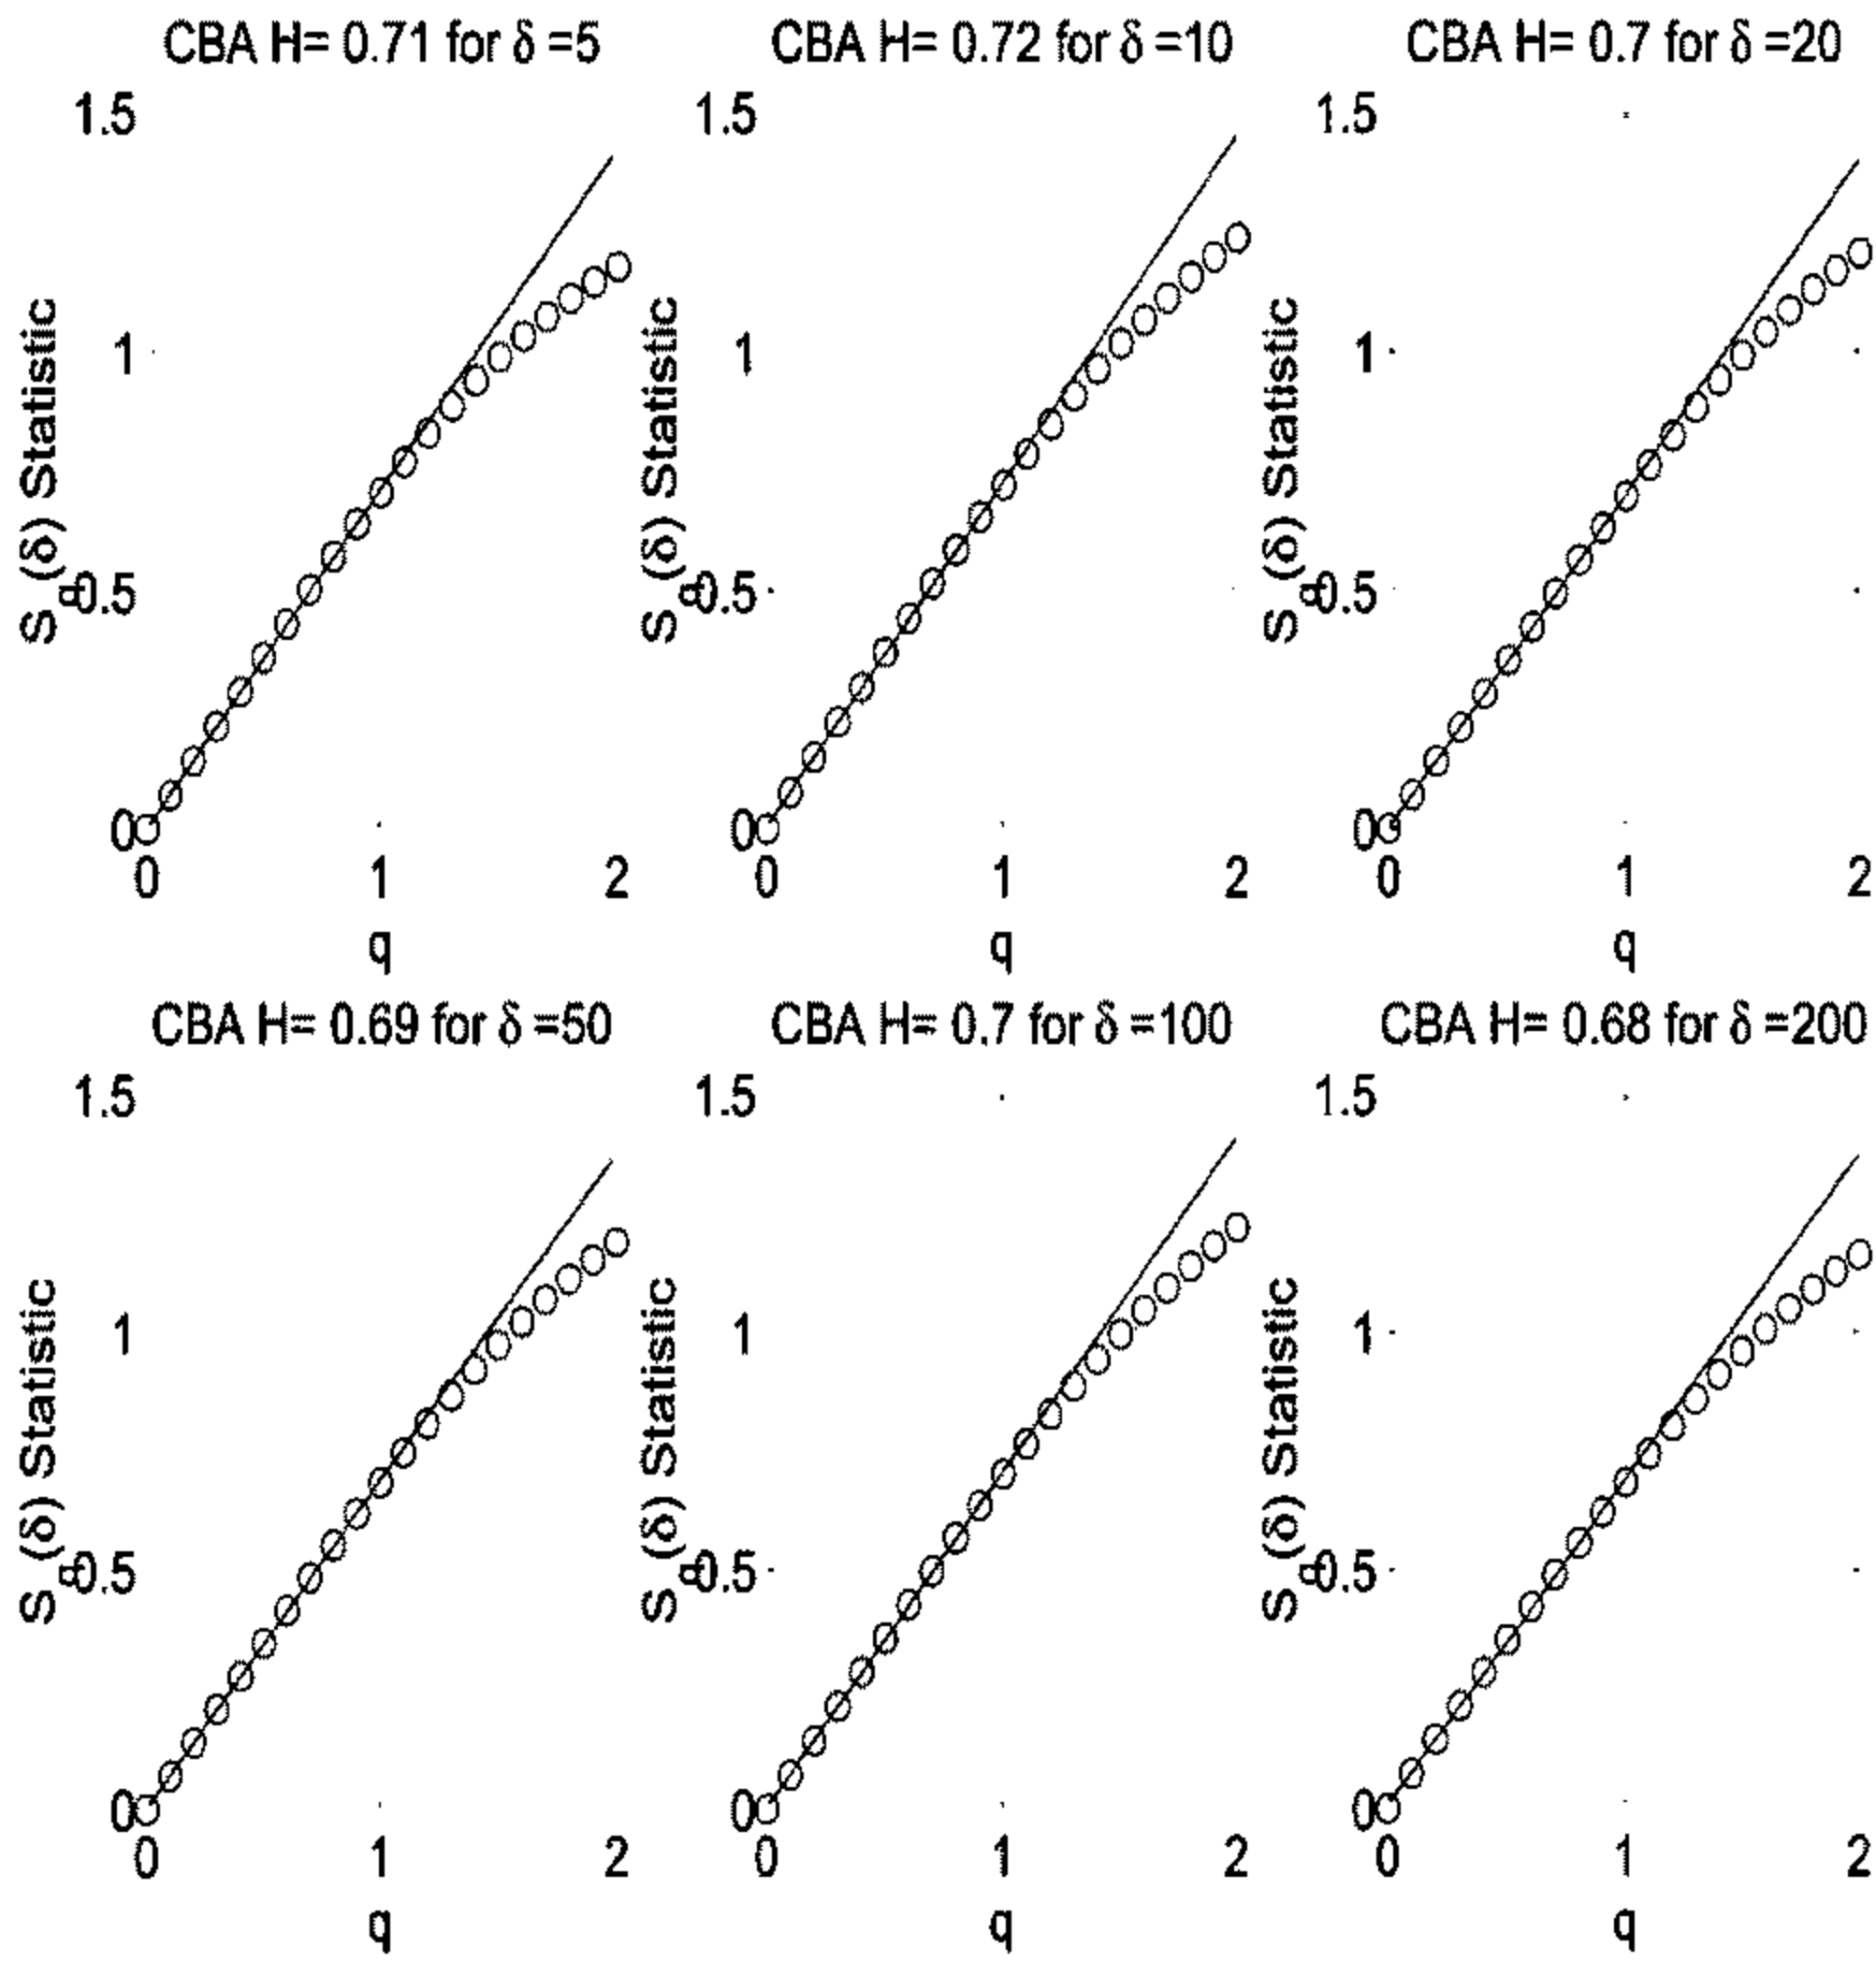


Figure 2.7: Plots of the statistic $\frac{\log \mathbb{E}[S_q(\delta)] - \log \mathbb{E}[S_q(1)]}{qH \log \delta}$ vs q . It is clear that the plots are not linear for $q > 1$ and thus do not support Heyde's assumptions.

are attainable. However, FATGBM also contains some drawbacks. Estimation of parameters is unclear and potentially difficult. Also it is not clear from examining the empirical data that the underlying assumption of $T_t - t$ self-similarity is correct.

2.4.3 Multifractal model of asset returns

Building on the earlier work of Mandelbrot (1963) and Mandelbrot and Taylor (1967) the multifractal model of asset returns (MMAR) was proposed by Mandelbrot, Calvet and Fisher (1997). The MMAR generalises Mandelbrot's previous models to include multifractal scaling.

A multifractal process $X(t)$ satisfies the relation,

$$\mathbb{E}[|X(t)|^q] = c(q)t^{\zeta(q)+1}. \quad (2.113)$$

Here $\zeta(q)$ is known as the scaling function of the multifractal process $X(t)$. Multifractal processes also satisfy the general scaling relation,

$$X(ct) \stackrel{d}{=} M(c)X(t) \quad (2.114)$$

where $M(c)$ is a stochastic process independent of $X(t)$ (see Reidi (2003)). Note that a self-similar process with index H is a multifractal with scaling function $\zeta(q) = Hq - 1$ since,

$$X(t) \stackrel{d}{=} t^H X(1) \quad (2.115)$$

$$\Rightarrow \mathbb{E}[|X|^q] \stackrel{d}{=} t^{Hq} \mathbb{E}[|X(1)|^q] \quad (2.116)$$

$$\stackrel{d}{=} c(q)t^{(Hq-1)+1} \quad (2.117)$$

where $c(q) = \mathbb{E}[|X(1)|^q]$ which is consistent with Eq. (2.113).

Under this model the asset price P_t is postulated to be a compound process of the form,

$$\log P_t = B_H(\Theta_t) \quad (2.118)$$

where $B_H(t)$ is a fractional Brownian motion with self similarity index H and Θ_t , called multifractal trading time, is a multifractal function of clock time. Specifically Θ_t is a multifractal process, independent of $B_H(t)$, with continuous non-decreasing

paths and stationary increments. The resulting process, $\log P_t$, is a multifractal with scaling function $\zeta_P(q) = \zeta_\Theta(Hq)$ and stationary increments.

The MMAR allows for a quite sophisticated dependence structure to be introduced into the model. The log-return itself may display correlations through the employment of a fractional Brownian motion, while the higher order correlations of the log-return (namely the absolute and squared return) inherit a dependence structure from the multifractal activity time. The MMAR also possesses general scaling behaviour and fat tailed distributions.

Fisher, Calvet and Mandelbrot (1997) and Calvet, Fisher and Mandelbrot (1997) explored the properties of the MMAR in detail and conducted an empirical study in order to ascertain the multifractal properties of real data. Their model is quite general and able to produce interesting behaviour, although it faces a myriad of problems from a practical point of view. The theory of multifractals is little understood and numerical methods remain somewhat inaccessible. Other than the binomial multifractal cascade, specific multifractal processes are few and far between. Also it is widely acknowledged that asset returns are uncorrelated and do not warrant a fractional Brownian motion. The FBM seems to be included just to make the MMAR as general as possible. Aside from these problems, the issue of multiscaling and multifractality is an interesting one which must be examined in greater detail. The analysis of data for multifractal properties requires high frequency data as long time series are needed. This idea will be explored further in the next chapter.

2.5 Jump-diffusion processes

Jump-diffusion processes are popular in modelling asset returns. The models were developed specifically to model large sudden price jumps. These models are typically constructed using a Poisson process. Press (1967) first asserted that the price process could be accounted for by aggregating a random number of price changes of random size and super-imposing a Brownian motion. Merton (1975) introduced the jump-diffusion model and examined the effects the model had on option pricing. Cox and Ross (1975) examined pure jump processes, their diffusion limits and option pricing

under alternative diffusion processes. Further work has recently been performed in this area by Madan (2002) in representing a variance gamma process as a pure jump process.

2.5.1 The Merton model

The critical assumption of Black and Scholes is that trading takes place continuously in time and the price dynamics of the stock price have continuous sample paths. In order to modify this assumption, Merton (1975) proposed the so called jump-diffusion process. Under this process the stock price path is discontinuous. This assumption was made in order to ascertain what would happen to the Black-Scholes model in the presence of non-local changes, or jumps, in price.

The jump-diffusion model accepts random price changes of two types. Normal shocks in price are caused by fluctuations in supply and demand, and other changes in information that cause marginal changes in stock value. This is what is modelled by GBM with constant variance and continuous sample paths. Conversely, abnormal shocks are due to the arrival of important new information about the asset that have more than a marginal affect on the price. The times at which such arrivals occur are random. This component is described by a jump process. Merton chose the Poisson process as a natural way to represent this jump process.

The proposed asset price is written as a stochastic differential equation,

$$dP_t = P_t[(\alpha - \lambda k)dt + \sigma dW_t + dq_t] \quad (2.119)$$

where q_t is a Poisson process independent of W_t and λ is the mean number of arrivals per unit time. The parameter $k = \mathbb{E}[(Y - 1)]$ is the drift associated with the Poisson process, where $(Y - 1)$ is the random percentage change in the stock price if the Poisson event occurs.

The solution of the above stochastic differential equation can be written,

$$P_t = P_0 e^{[(\alpha - \sigma^2/2 - \lambda k)t + \sigma W_t]Y(n)} \quad (2.120)$$

with $Y(0) = 1$; $Y(n) = \prod_{j=1}^n Y_j$ where Y_j are independent identically distributed variables and n is a Poisson distributed random variable with parameter λt . Merton

(1973) studied the case where the Y_j are drawn from the lognormal distribution. In this case the model for P_t becomes lognormal with variance given by a Poisson random variable, similar to the mixture model of Press (1967).

Merton concentrated on pricing options under this framework. He emphasised that although the introduction of jumps seems to explain the divergence of Black-Scholes from the real market it would be presumptuous to claim that this was the true reason since other deviations from the Black-Scholes assumptions may also explain this divergence. While this model admits some analytic tractability, implementation can prove difficult as in practice it is difficult to estimate the jump process parameters, especially the jump frequency, λ .

2.5.2 The Cox-Ross model

Cox and Ross (1975) examined the use of pure-jump models in option pricing. They presented two Poisson based models. The first is an example of a simple pure birth and death process,

$$dP_t = \begin{cases} k^+ - 1 & \text{with probability } (\pi^+ \lambda P_t dt) \\ k^- - 1 & \text{with probability } (\pi^- \lambda P_t dt) \\ 0 & \text{with probability } (1 - \lambda P_t dt) \end{cases} \quad (2.121)$$

where $k^+ - 1 > 0$ and $k^- - 1 < 0$ are the jump amplitudes. Here, $\lambda P_t dt$ represents the probability that an event occurs in time dt . An event is either the ‘birth’ of $k^+ - 1$ new units (with probability π^+) or the ‘death’ of $1 - k^-$ old units. We see here that the intensity, $P_t \lambda$, is proportional to asset value. Passing to the diffusion limit via the Kolmogorov backward equation, Cox and Ross write the limit of Eq. (2.121) as

$$dP_t = \mu P_t dt + \sigma \sqrt{P_t} dW_t. \quad (2.122)$$

The process described by this stochastic differential equation is known as the square root process.

The second case explored occurs when the process contains a drift term and the

intensity, λ , is constant. This yields,

$$dP_t = \mu P_t dt + \begin{cases} k^+ - 1 & \text{with probability } (\pi^+ \lambda dt) \\ k^- - 1 & \text{with probability } (\pi^- \lambda dt) \\ 0 & \text{with probability } (1 - \lambda P_t dt). \end{cases} \quad (2.123)$$

Taking the limit of this process they arrived at the stochastic differential equation,

$$dP_t = \mu P_t dt + \sigma dW_t. \quad (2.124)$$

This process is commonly known as the Ornstein-Uhlenbeck process and is used in statistical physics to describe a particle moving in a parabolic potential in the presence of external thermal noise.

Cox and Ross (1975) used these two stochastic processes to model the underlying asset. These models are special cases of what has come to be known as the constant elasticity of variance (CEV) model which has the general form,

$$dP_t = \mu P_t dt + \sigma P_t^\delta dW_t. \quad (2.125)$$

Although these two processes have not had much impact on the modelling of the asset price, both the square root and Ornstein-Uhlenbeck process have been the focus of stochastic volatility in the theory of option pricing.

2.6 Stochastic volatility

Stochastic volatility models were developed to counter deficiencies of the Black-Scholes option pricing model. The assumption of a constant volatility parameter is recognised as a weakness in the Black-Scholes model by both researchers and practitioners. It is now widely accepted that the constant volatility of the asset price under GBM is insufficient to describe real data, as we have seen in the previous sections. After the market crash in 1987 it became clearer that the Black-Scholes model had marked discrepancies when compared to actual fair prices of contracts. These discrepancies gave rise to the famed volatility smile and term structure of volatility. In the same year of the market crash the first of a new generation of option pricing

models appeared and stochastic volatility models were born. These models seek to explain the discrepancies between Black-Scholes and real world prices through the introduction of a stochastic volatility parameter. In this sense, stochastic volatility models are similar to the subordinated processes of the previous two sections. They offer a slight advantage over the early subordinated models in that the volatility may be correlated to the underlying Wiener process. However, these models are usually proposed ad hoc, without physical reasoning.

In a continuous time framework the volatility, V_t is assumed to follow a diffusion process. The stochastic volatility model has the general form

$$dP_t = \alpha_t P_t dt + \sigma_t P_t dW_t^1 \quad (2.126)$$

$$dV_t = b(V_t)dt + \xi\eta(V_t)dW_t^2 \quad (2.127)$$

where W_t^1 and W_t^2 are Wiener processes with correlation ρ . In this equation α_t is the drift rate and σ_t is the instantaneous volatility given by, $V_t = \sigma_t^2$. The parameter ξ is the ‘volatility of volatility’ whilst the functions $b(V)$ and $a(V) = \xi\eta(V)$ are the drift and diffusion coefficients of the volatility process. Many different processes have been proposed under stochastic volatility. In practice, stochastic volatility models are an important part of financial mathematics particularly in relation to option pricing. However, from an asset dynamics point of view these models have been proposed with the desire to create analytic and usable pricing models, rather than as an attempt to explain empirically observed phenomena. Consequently in this section we will limit our investigation to the main stochastic volatility models, restraining ourselves from analysing their distributions and other stochastic properties.

2.6.1 The Hull and White Model

In the same year as the Black Monday financial crash, Hull and White (1987) generalised the Black-Scholes model to allow for stochastic volatility. They took the volatility to be described by a geometric Brownian motion,

$$dV_t = \alpha V_t dt + \xi V_t dW_t. \quad (2.128)$$

With this assumption the volatility displays lognormal increments similar to the Clark (1973) subordinated model. The Hull and White model does not allow for

closed form solutions for option pricing and requires numerical methods to solve the resulting partial differential equation.

2.6.2 The Ornstein-Uhlenbeck Model

Stein and Stein (1991) assumed that the volatility follows an Ornstein-Uhlenbeck process,

$$dV_t = \beta V_t dt + \xi dW_t. \quad (2.129)$$

This is a mean reversion process and is stationary, unlike the Hull and White model. This model is popular as the OU-process is a well defined stochastic process with a great deal of literature devoted to it. Assuming further that the volatility is uncorrelated to the log-return (*i.e.* $\rho = 0$) option prices are obtained through an average of Black-Scholes formula values over different volatility paths.

2.6.3 The square root model

The square root model was originally developed by William Feller and applied to option pricing by Heston (1993). This process can be obtained from the Ornstein-Uhlenbeck process. If the volatility follows,

$$d\sqrt{V_t} = \beta\sqrt{V_t}dt + \xi dW_t \quad (2.130)$$

then by Itô's formula it can be shown that the variance, V_t satisfies the process,

$$dV_t = (\omega - \theta V_t)dt + \xi\sqrt{V_t}dW_t. \quad (2.131)$$

This model is an important case of stochastic volatility. Heston (1993) derived semi-closed form solutions for option prices under this model. These extend to the case where the volatility is correlated to the asset price (*i.e.* $\rho \neq 0$), allowing for skewness effects in the volatility surface.

2.6.4 The 3/2 model

The 3/2 model (Lewis (2000)) is described by the stochastic differential equation

$$dV_t = (\omega - \theta V_t)dt + \xi V_t^{3/2} dW_t. \quad (2.132)$$

This model is important because it exhibits closed form solutions for option prices and displays features in addition to those of the square root model.

2.7 GARCH and generalised ARMA models

In this section we look at the discrete time series approach to asset modelling. By far the most popular of the models in this class are the renowned ARCH and GARCH models of Engle (1982) and Bollerslev (1986). The ARCH and GARCH models were developed from the discrete ARMA time series models of Box and Jenkins (1970). Other extensions of these models include ARIMA and fractional ARIMA time series. We also include in this section the generalised ARMA model of Peiris (2004), describing a possible misclassification problem that can arise when using time series models of this type.

2.7.1 GARCH

The GARCH (generalised auto-regressive conditional heteroscedasity) model was developed by Bollerslev (1986) from the ARCH model of Engle (1982). The huge impact that this model has had on the world of finance can be seen through the awarding of the 2003 Nobel Prize in Economics to Robert Engle and Clive Granger for their subsequent work on this and similar discrete time series models.

GARCH is a simple model that is able to describe time series that while asymptotically stationary, can display local non-stationary behaviour. The GARCH(p, q) model is defined as follows,

$$X_t = \sigma_t \epsilon_t \quad (2.133)$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^p \alpha_i X_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j}^2 \quad (2.134)$$

here ϵ_t is a random variable usually drawn from the normal distribution, $\epsilon_t \sim N(0, 1)$. We can see that the variance is conditional on past values through the second summation term appearing in Eq. (2.134). GARCH has the property that while the

residuals themselves are uncorrelated, the squared residuals, X_t^2 are correlated with each other.

The GARCH model that is most often used in practice is GARCH(1,1). We will examine this process as general results for GARCH(p, q) are not analytically obtainable. GARCH(1,1) is specified by the difference equation,

$$\sigma_t^2 = \alpha_0 + \alpha X_{t-1}^2 + \beta \sigma_{t-1}^2. \tag{2.135}$$

The correlation between the squared residuals, X_t^2 and X_{t+k}^2 is given by Ding and Granger (1996),

$$\rho(k) = \text{corr}[X_{t+k}^2, X_t^2] = \left(\alpha + \frac{\alpha^2 \beta}{1 - 2\alpha\beta - \beta^2} \right) (\alpha + \beta)^{k-1} \tag{2.136}$$

provided that the model has a finite fourth moment, which in turn implies $0 < \alpha < 1/\sqrt{3}$. The autocorrelation function for the GARCH(1,1) model decreases exponentially (Granger and Ding (1996)).

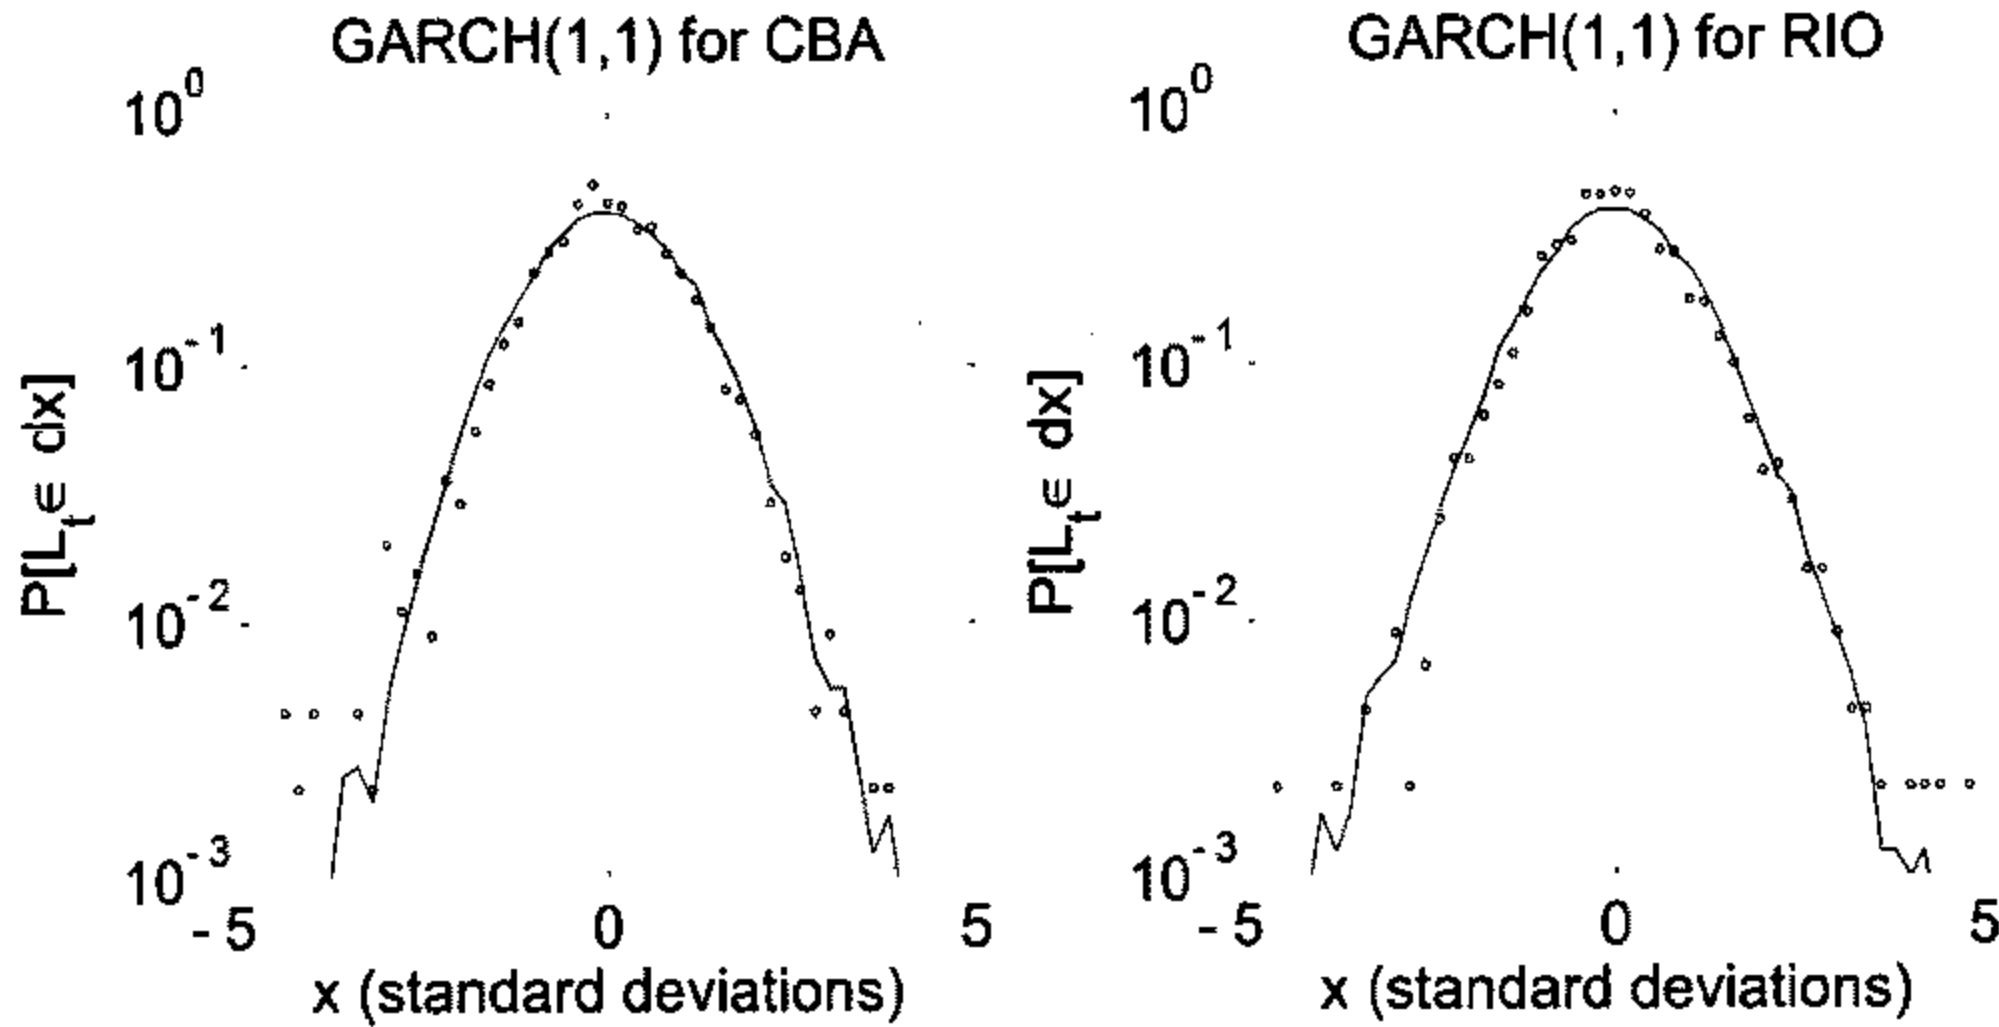


Figure 2.8: The GARCH(1,1) model fits the empirical density function of the daily log-return quite well.

In Fig. 2.8 we have fitted GARCH(1,1) to ASX daily return data using maximum likelihood estimation. We see from Fig. 2.8 that GARCH(1,1) is able to produce a good fit to the log-return distribution. Fig. 2.9 displays the correlations

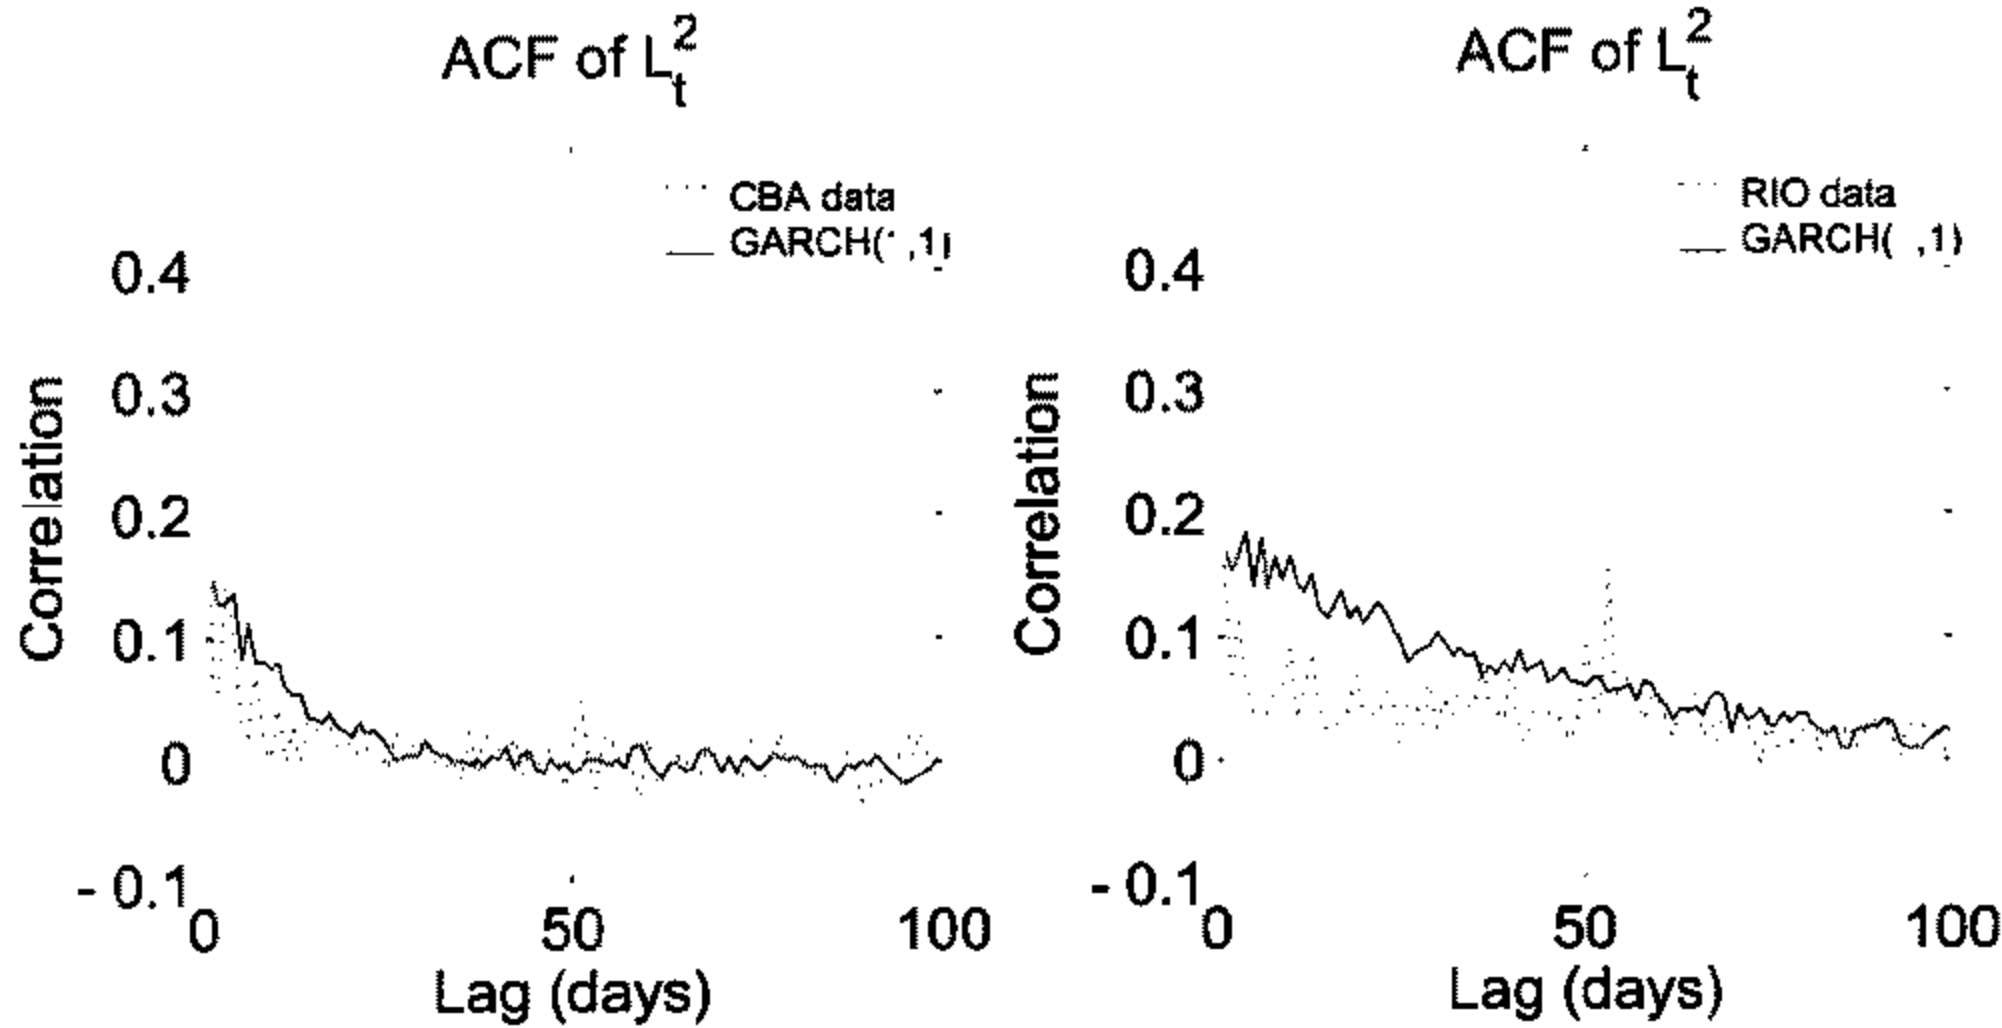


Figure 2.9: It can be seen in this plot that the autocorrelation structure of GARCH(1,1) does not fully describe that observed in real data.

of GARCH(1,1) compared to the real data. However, not much is known about the scaling behaviour of the GARCH model and in general the results may not be consistent over different time scales.

Much work continues to be done on the GARCH model mainly due to the fact that it is such a useful tool for practitioners. GARCH based option pricing models have been developed by Duan (1995), Kallsen and Taqqu (1998) and Heston and Nandi (1997). The model has sprouted many offshoots *e.g.* I-GARCH, H-GARCH, P-GARCH, each designed to incorporate different aspects of price behaviour.

2.7.2 Generalised ARMA models and misclassification

Autoregressive models, in particular AR(1), can be used as a discrete analogue for stochastic volatility (Musielka and Rutkowski (1999)). Here we present the work of Bertram and Peiris (2004) which examines a misclassification problem that can arise when modelling data which contains low frequency components. We illustrate this problem by fitting GAR(1) and AR(1) models to volume and transaction frequency data from the Australian Stock Exchange.

The modelling of time series which contain low frequency components is a prob-

lem of interest in many fields of study. In particular, low frequency events in financial time series often represent the most important events, *i.e.* the booms and crashes of the market. An important case of these low frequency components is in the study of long memory processes, see Beran (1994). The motivation for this work arose from the paper by Granger and Teräsvirta (1999) which discussed a possible misclassification between nonlinear models and $I(d)$ processes. In Peiris (2004) it was shown that modelling a time series of varying frequency components with an ARMA process leads to a misclassification problem. In this section we show that modelling such a time series with a GAR(1) or a AR(1) model yields similar results for the regression parameter, however for the case of the GAR(1) model we have another parameter that is able to describe the degree of frequency of the data. This result indicates how the ARMA model is unable to distinguish between different time series of this type.

Let B be the backshift operator and I be the identity operator such that $B^j X_t = X_{t-j}$ and $I = B^0$. The model we consider here is the generalised autoregressive of order 1 (GAR(1)) process given by

$$(I - \alpha B)^\delta X_t = Z_t, \quad (2.137)$$

where $|\alpha| < 1$, $\delta > 0$ are constants and $\{Z_t\}$ is a sequence of uncorrelated random variables with mean 0 and variance σ^2 (see Peiris et. al. (2003)). When $\delta = 1$ the model reduces to an AR(1) process.

The spectral density of the GAR(1) model is given by

$$S_X(\omega) = \frac{\sigma^2}{2\pi} (1 - 2\alpha \cos \omega + \alpha^2)^{-\delta}; \quad -\pi \leq \omega \leq \pi. \quad (2.138)$$

Thus it can be seen how the parameter δ controls the degree of frequency of the process. For example, for $\delta > 0$ it is clear that $S_X(\omega)$ is large in neighbourhoods of $(\omega, \alpha) = (0, 1)$ and $(\omega, \alpha) = (\pi, -1)$. The autocovariance of this process has been shown in Peiris (2004) to be given by

$$\gamma_0 = \text{Var}(X_t) = \sigma^2 F(\delta, \delta; 1; \alpha^2) \quad (2.139)$$

$$\gamma_k = \text{Cov}(X_t, X_{t+k}) = \frac{\sigma^2 \alpha^k \Gamma(k + \delta) F(\delta, k + \delta; k + 1; \alpha^2)}{\Gamma(\delta) \Gamma(k + 1)}; \quad k > 0 \quad (2.140)$$

where $\Gamma(x)$ is the gamma function and $F(a, b; c; d)$ is the Gauss hypergeometric function, given in Abramowitz and Stegun (1968). Examples shown in Peiris (2003) show that the autocorrelation (ACF), partial autocorrelation (PACF) and spectral density are similar for some values of α and δ .

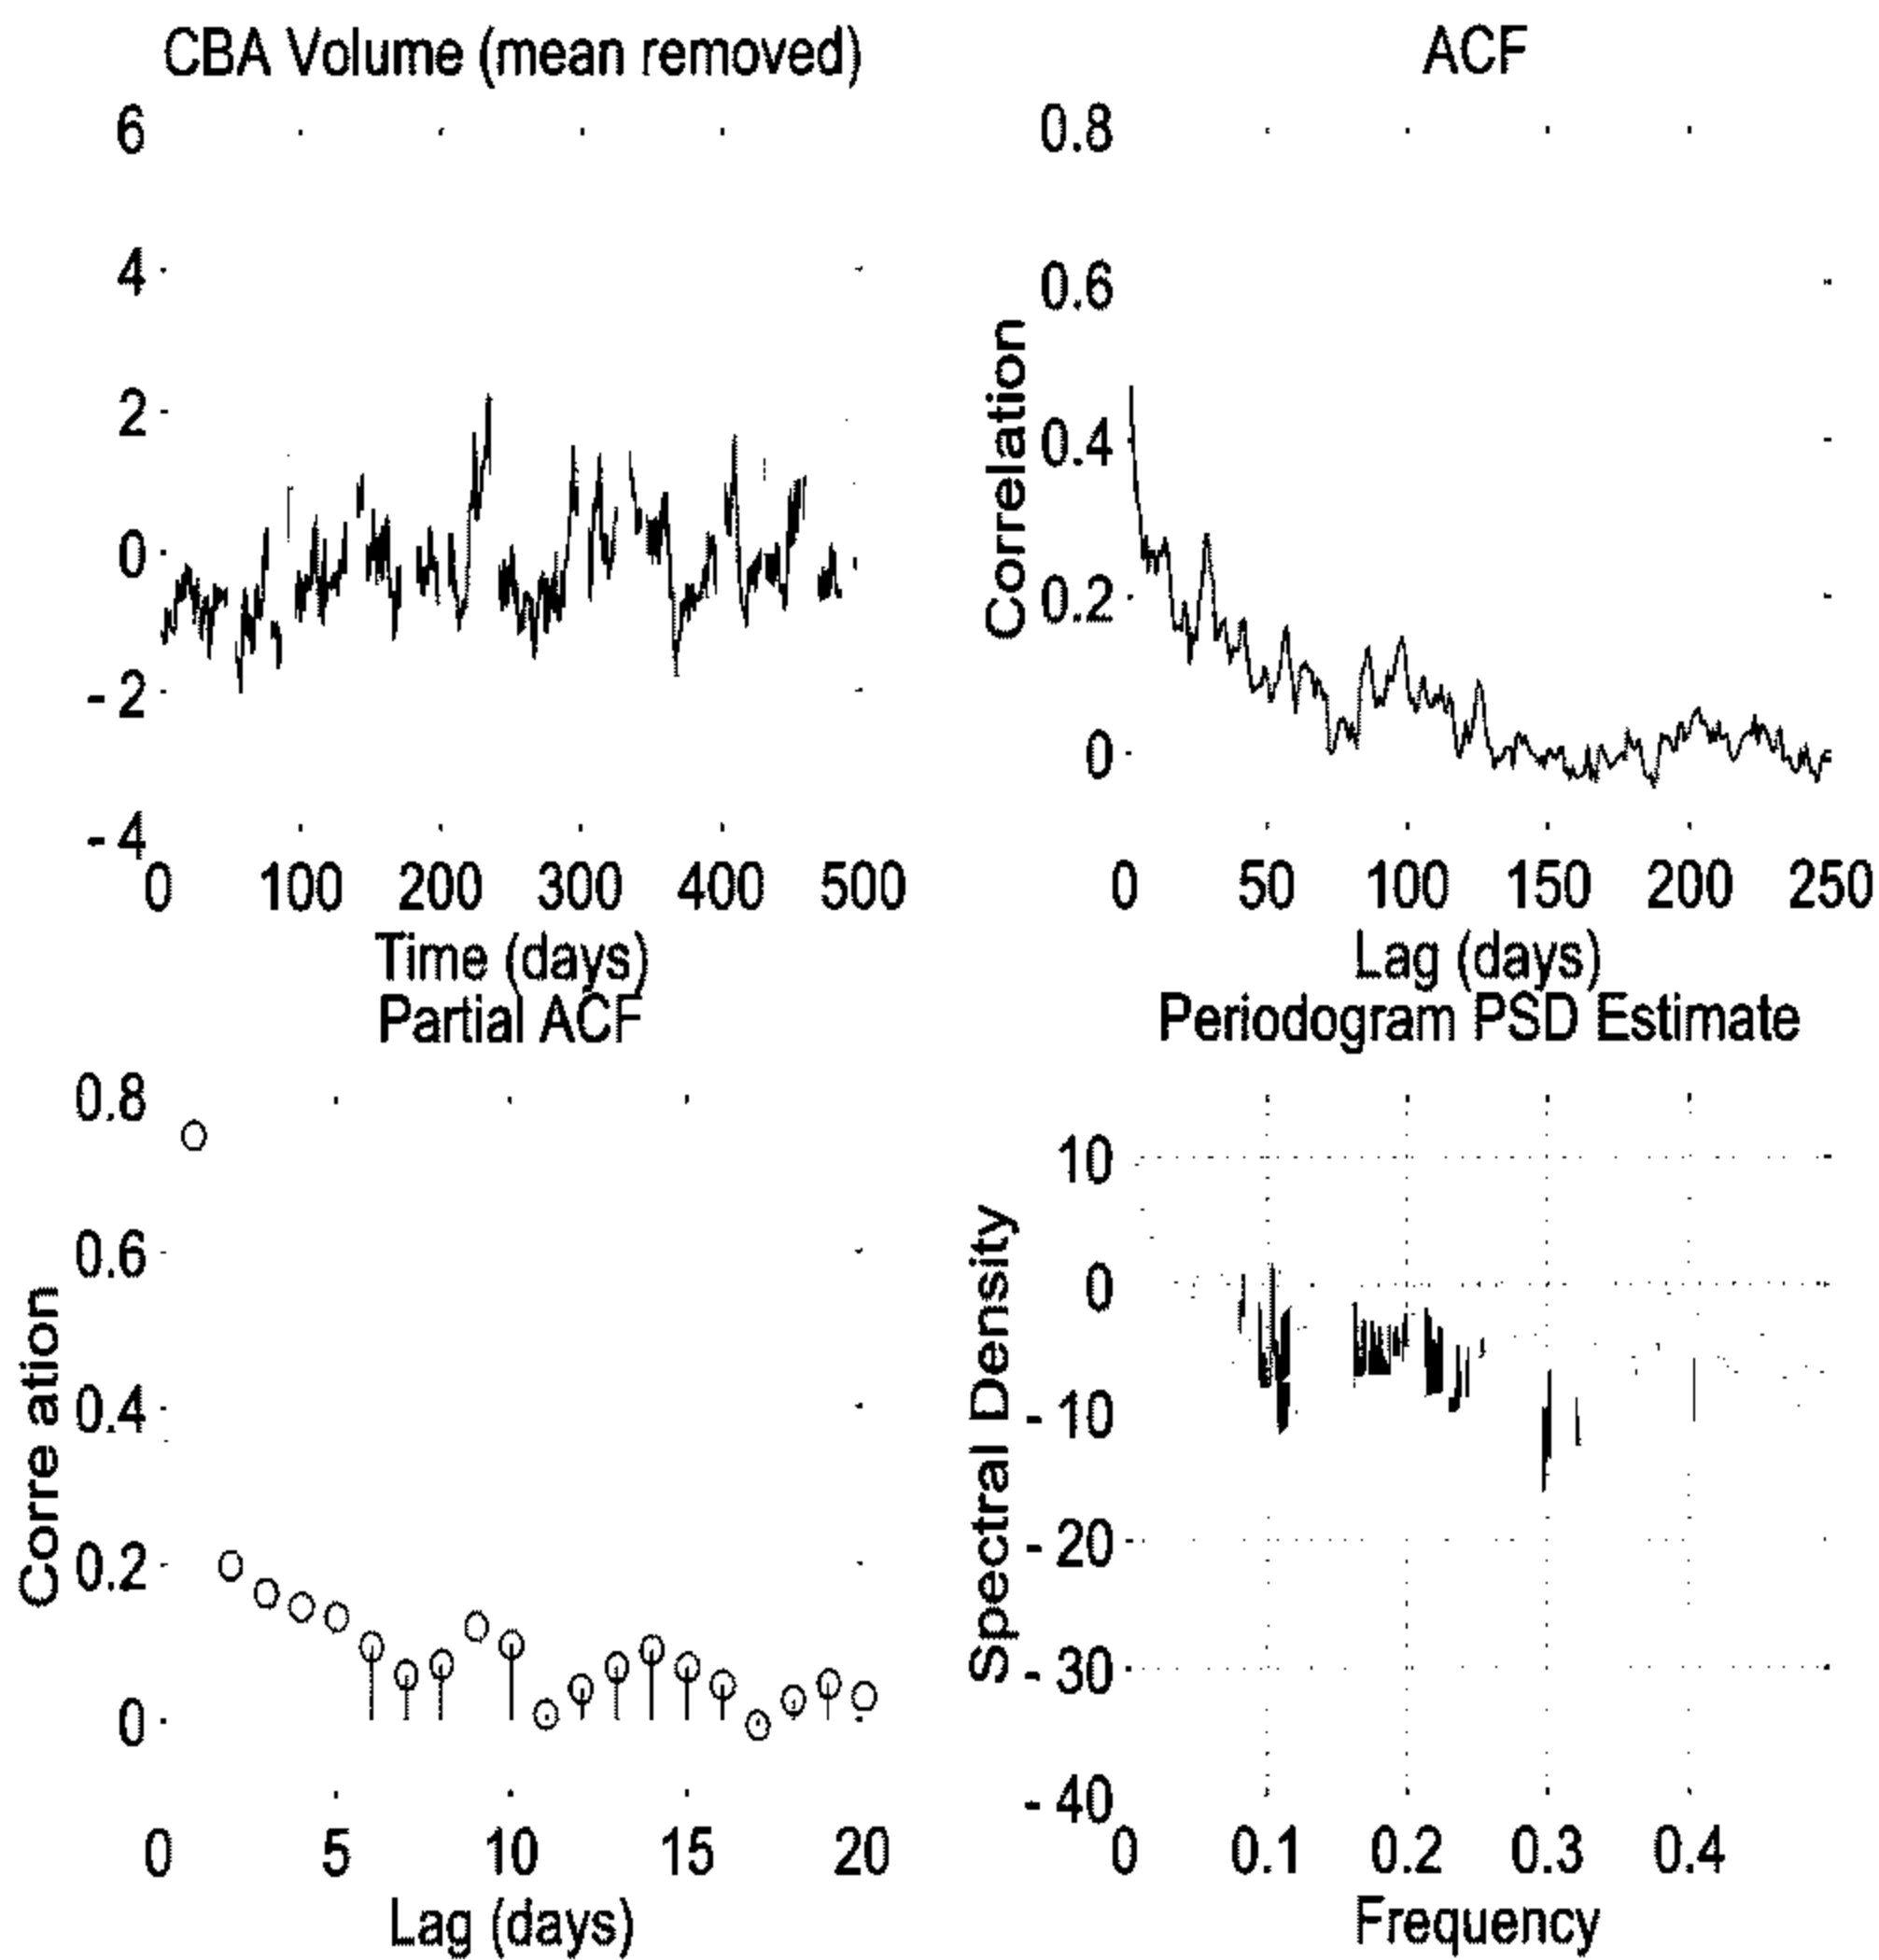


Figure 2.10: The plots show the time series ACF , PACF and spectral density.

We fit both GAR(1) and AR(1) models to Australian Stock Exchange data in order to demonstrate the misclassification problem that arises when modelling data of varying frequency. We use daily traded volume and transaction frequency

data, both of which have been proposed as measures for the stock price volatility (Bertram (2004)). The ACF, PACF and spectrum shown in Fig. 2.10 indicate that an AR(1) model may be suitable to describe the data. We emphasise that we are not suggesting GAR(1) or AR(1) as a plausible model for this process, rather using the data as an example of the problem encountered when dealing with a time series of this type. Empirical studies of this data have shown it to possess slowly decaying correlations and hence a fractionally differenced ARIMA(1, d ,0); $d \in (0, 0.5)$ would be an alternative model. It must be noted that the time series for the volume is bound below by zero, however since the mean of the process is quite large ($\mu \sim 10^6$) we ignore the effects that this lower bound has on the time series. Table A.8 displays the estimated GAR(1) and AR(1) parameters for both the volume and transaction frequency, calculated for a selection of stocks. Looking at the AR(1) parameters for volume we find that they fall roughly into two categories. The stocks CML, RIO, WMC suggest a regression parameter $\alpha_{AR} \sim 0.50 \pm 0.1$ while the stocks BHP, BIL, CBA indicate $\alpha_{AR} \sim 0.80 \pm 0.1$. Examining the GAR(1) parameter values we find that the regression parameter is approximately the same as that of the AR(1) model, $\alpha_{GAR} \approx \alpha_{AR}$. However, within the two categories we find that the value of δ varies greatly. This result indicates the inability of the ARMA models to discriminate between the different time series. For instance, Table A.8 suggests the same AR(1) model for both WMC and RIO, while the GAR(1) parameters show the extent of their differing frequency components. Turning our attention to the transaction frequency we find results similar to those of the volume. For most of the stocks, the parameter estimates indicate an AR(1) model with $\alpha \sim 0.75 \pm 0.1$. Looking at the GAR(1) estimates we find that although the regression parameter estimates are roughly the same, the frequency parameter values vary wildly, with δ values lying in the range (0.5, 1).

A forecast on standardised CBA volume (mean removed) was performed with both AR(1) and GAR(1) models. The three forecasts from the origin at $t = 1000$, 1001 and 1002 days are displayed in Table A.9. From these results we can see that for the first two GAR(1) forecasts are closer to the observed values than those of AR(1).

We also see here how the extra parameter of the GAR(1) model is able to control the degree of frequency of the model and thereby describe data that consists of varying frequency components. The GAR(1) and AR(1) model parameters were estimated from volume and transaction data from a selection of equities traded on the Australian Stock Exchange. It is well known that the data used in this study contain low frequency components that lead to a slowly decaying autocorrelation function.

2.8 Summary

This chapter has provided a review of previous models developed for price dynamics over the past one hundred years. We examined the pioneering work of Bachelier and the development of his work that lead to geometric Brownian motion. The introduction of the stable Paretian hypothesis by Mandelbrot was a critical point for financial modelling. Despite the debate that continues to rage over the assumption of infinite variance, Mandelbrot's work sparked an interest in developing alternative models for asset dynamics.

The subordinated processes, first used by Mandelbrot and Taylor (1967) and Clark (1973), provided a new class of asset models from which a plethora of models has emerged. We examined four of the main subordinated models and performed parametric estimation for each model, including GBM, on Australian equity daily log-return data. Each model's ability to describe the data was measured via the log-likelihood function. We found that the Student- t model of Praetz (1972) best described the data with $\nu = 4$, indicating a finite variance. The Madan and Seneta (1990) variance gamma model, the Clark (1973) lognormal, and the Mandelbrot and Taylor (1967) Lévy-stable model all followed closely behind, with all the models out performing GBM. The results obtained from this estimation showed how each model's description of distribution tail behaviour differed. Due to the low frequency nature of the data set, it is impossible to distinguish the tail behaviour of daily data. This is precisely the problem that has faced researchers in the past and caused so much speculation as to the true behaviour of the stochastic process that drives price

change.

It has been shown in studies such as Granger and Ding (1996) and Ding and Granger (1996) that price data indicate the presence of higher order correlations in the log-return. We showed an example of CBA data that appear to display an uncorrelated log-return, but with a distinct correlation present in the absolute and squared log-return. Moving on from the classic subordinated models, we examined three more recent models for price change that extended the class of subordinated processes to include such dependence features. The assumptions of Heyde (1999) FATGBM model were tested by empirically constructing the Activity Time, T_t , and examining $T_t - t$ for fractal behaviour. Our results from ASX daily data were not in complete agreement with Heyde's assumptions.

Time-varying volatility models and stochastic volatility models represent important developments in the history of price dynamics. Both of these model types are central to modern option pricing. The discrete time series model of GARCH presents one of the most important asset modelling tools for time series analysis. We examined the properties of the GARCH model and fitted the model to daily log-return data. In the last section of the chapter we examined the problems that can arise in discrete time series modelling. We showed how a non-invertibility problem can occur whereby the ARMA model is unable to represent the different time series. The corresponding GARCH(1) model was able to distinguish between the different time series. The same auto-regression parameter of the AR(1) model, in addition to a frequency parameter, δ , to describe the degree of frequency of the data. The GARCH(1) parameter GARCH(1) is able to distinguish between the time series of different stocks based on their various degrees of frequency without significantly changing other statistical properties of the data (*e.g.* ACF, PACF). Note that the fractional ARIMA models can represent only two frequency components, whereas GARCH models are able to describe varying frequency components.

While we have not set out specifically to model the time series in question with GARCH(1), we have used the presented examples to highlight a problem that has strong repercussions in forecasting. The use of a standard ARMA type model will be insufficient when forecasting time series containing low frequency components.

The inability of the standard ARMA model to represent the differing frequency properties of this type of data is a major drawback to its functionality. The use of the GAR model prevents the misclassification problem occurring and should lead to more accurate forecasts.

Chapter 3

Tests and analysis of ASX equity

Data

As we discussed in Chapter 1, geometric Brownian motion has long been criticised as offering a poor description of stock prices. Over the years numerous improvements to geometric Brownian motion have been put forward and many other models have been proposed. However these improved models were usually based on speculated and assumed behaviour of financial markets rather than on observed behaviour. The reason for this approach was mainly due to a lack of financial data on which to build a knowledge base about the real behaviour of financial markets. Daily returns were the most common quantity from which to build a picture of the market behaviour. But as we have seen in Chapter 2 such time series generate inconclusive results as to the nature of stock price dynamics.

In this chapter we turn our attention to the use of high frequency financial data³. We examine the ability of previous models to describe the high frequency log-return distribution and present a comprehensive empirical study of ASX equity data. We present the fundamental properties of the distribution, the dependence structure, and the scaling properties of the data. In the past many models were constructed around a mathematical property chosen to give the model appealing theoretical features. The results and observations obtained from this study will allow for the

³The results in this chapter are based on the top 200 ASX equities ranked by turnover

development of models that are data driven as opposed to axiomatically driven. Much of what is presented in this chapter appears in published form in Bertram (2004) while the remainder of the work extends the contents of that publication.

3.1 Econophysics and empirical finance

During the 1980's a revolution of sorts was taking place in financial markets around the world. This revolution consisted of the transition to complete electronic stock exchange operation. Exchanges switched from the traditional system of trading pits, with physical interaction between traders and transactions recorded by hand, to a situation where market participants conducted business through a central computer network. This change has allowed for exchanges to keep track of virtually all events that take place in the market from the release of company reports, to individual bid and ask quotes on a stock. The rise of modern computing power, combined with advances in data storage and the advent of the internet have resulted in the creation of huge databases of financial information. These databases are run by companies primarily as valuable source information to practitioners. However an increasing number of companies are making data available for research.

The term high frequency data refers to financial data that has a typical sampling period of less than a day. A sampling period ranging anywhere from individual transactions (tick-by-tick) to several hours in length will produce a time series suitable of the name high frequency. Before high frequency data was available, researchers usually had to content themselves with daily or even weekly data. It was on this type of data that many of the early theories and models were tested. The use of such low frequency data results in a relatively short time series (*i.e.* 10 years of daily prices yields only ~ 2500 entries). The results obtained with such short time series can be inconclusive, as the sample will be too small to accurately summarise the behaviour of the underlying stochastic process. For example, while studying the distribution properties of financial data we are especially interested in large standard deviation (*i.e.* rare) events. When fitting the subordinated models in Chapter 2 we saw how the distribution tails could not be properly represented by low frequency daily data

and thus would not accurately represent the distribution of large standard deviation events. High frequency data allows us to examine extremely long time series (to the order of 10000 to 1000000 entries) and permits studies that were previously impossible. With high frequency data, financial markets have become an interesting area of study for those interested in the behaviour of complex systems, particularly statistical physicists.

Financial markets are open systems that have many interacting parties operating in the presence of feedback. As pointed out by Mantegna and Stanley (2000), these markets are continuously monitored and operate under well defined, stable governing rules that dictate the operation of the exchange. It is for these reasons that many physicists interested in the description, classification and modelling of complex systems have taken an interest in finance.

The influx of physicists, including the pioneering work of H. Eugene Stanley, has led to the creation of the field of econophysics, the science of empirical finance. While economists and mathematicians have traditionally been seen as the investigators of financial systems, econophysics differs from these approaches by way of the emphasis that it places on the empirical analysis of data. The idea is to move away from the strong theoretical and axiomatic approaches of financial mathematics and economics and take a less constructive, more descriptive approach. Econophysics aims towards the complete statistical characterisation of the stochastic process of price change. The main properties being the shape of the distribution of price change, temporal memory, scaling and other higher order statistical properties. With this knowledge the ultimate goal is to develop a theoretical model that is able to encompass all essential features of real financial markets.

Econophysics is analogous to the fields of astronomy, geophysics and astrophysics in the sense that it is impossible to perform an experiment to prove true a given hypothesis. One is only able to confirm or reject theorems based on currently available data sets. This means that there is a fundamental need for studies on many different data sets as these all contribute to the general understanding of financial markets.

The aim in the following sections is to use Australian Stock Exchange equity

data to gain an understanding of the type of behaviour observed in the ASX. The investigation we undertake is invaluable as a source of information on the fundamental properties of high frequency financial data. The knowledge gained from this study will be used in subsequent chapters to develop and test models that are more representative of the true nature of equity markets and price dynamics.

3.2 ASX equity data

In this thesis we make use of data from the Australian Stock exchange. In Australia, financial researchers are at somewhat of an advantage in terms of access to good quality high frequency data. The Securities Industry Research Centre of Asia-Pacific (SIRCA) has partnered with several research institutions to provide both ASX and Reuters financial data to researchers interested in this area. It is from SIRCA that we obtained the data set used in this study.

One benefit of using this data set is that the ASX represents a smaller less liquid market than the US, UK and some European markets. This means that it is more reactive to participant behaviour, allowing for the observation of specific events that would otherwise be hidden on larger markets. The data set contains a record of every transaction that took place over the period January 1993 to July 2002, for each of the 200 most actively traded ⁴ stocks on the ASX. Owing to the comprehensive nature of the information stored, our data set is very large (3.5 gigabytes). In order to investigate the behaviour of the data properly, it is necessary to create subsets to study. We will now define these subsets. If P_t represents the price of a stock at time t then we define the log-return⁵ over interval size τ as

$$L_{t,\tau} = \log(P_{t+\tau}) - \log(P_t). \quad (3.1)$$

We also define $V_{t,\tau}$ to be the total share volume traded in the interval $[t, t + \tau]$ and $N_{t,\tau}$ to be the total number of transactions (transaction frequency) that took place in the interval $[t, t + \tau]$. These subsets of data are created from the original data set

⁴Stocks were ranked by total turnover during the studied period.

⁵The data have been corrected for dividends and other dilution effects.

using the MATLAB functions listed in Appendix D. These functions were written specifically by the author to facilitate easy manipulation of the data. By creating time series for different interval sizes, τ , we are able to investigate properties of the data ranging over all time scales. The resulting time series for the log-return, $L_{t,\tau}$, volume $V_{t,\tau}$, and transaction frequency $N_{t,\tau}$ spanning the entire sampled period with sampling frequency $\tau = 10$ minutes comprises some 85000 data points for each.

3.2.1 Overnight jumps and intraday noise

The ASX operates using the Stock Exchange Automated Trading System (SEATS). SEATS is an electronic order book that trades continuously between the hours of 10:00 and 16:00 Monday to Friday. Before the opening and after the close, the market enters a special 'pre-open' mode, where orders may be entered, adjusted and cancelled but not executed until a fixed time determined by the ASX. In the morning this starts at 07:00, giving three hours for traders to adjust trades to compensate for overnight information flow. A similar procedure occurs after the close at 16:00 until a random time between 16:05 and 16:06, when a single market auction takes place to clear the order book and set the official closing price. This procedure is performed to reduce both the end of day volatility and the possibility of market manipulation by large market participants. The ASX allows for after hours trades between brokers, requiring that they report their activities to the ASX. Further, the exchange contains several dual listed stocks trading on exchanges such as London or New York. These ASX operations clearly impact on the price forming process over non-trading hours. This is significant as it affects the way we analyse and model our data. In Fig. 3.1 we show a typical month of trading for Rio Tinto sampled every 10 minutes. The price shows large changes occurring during the overnight interval as indicated by the dashed lines. If we compare the time series for the 10 minute log-return over the whole period 1993-2002, as seen in Fig. 3.2, with the log-return where the overnight price jumps have been removed, we immediately see the impact that the overnight log-return has on the series as a whole and we can speculate about the large effect it must have on a stock's volatility. This leads to the interesting notion of treating the stock price as two separate processes: an intraday

process and an overnight process.

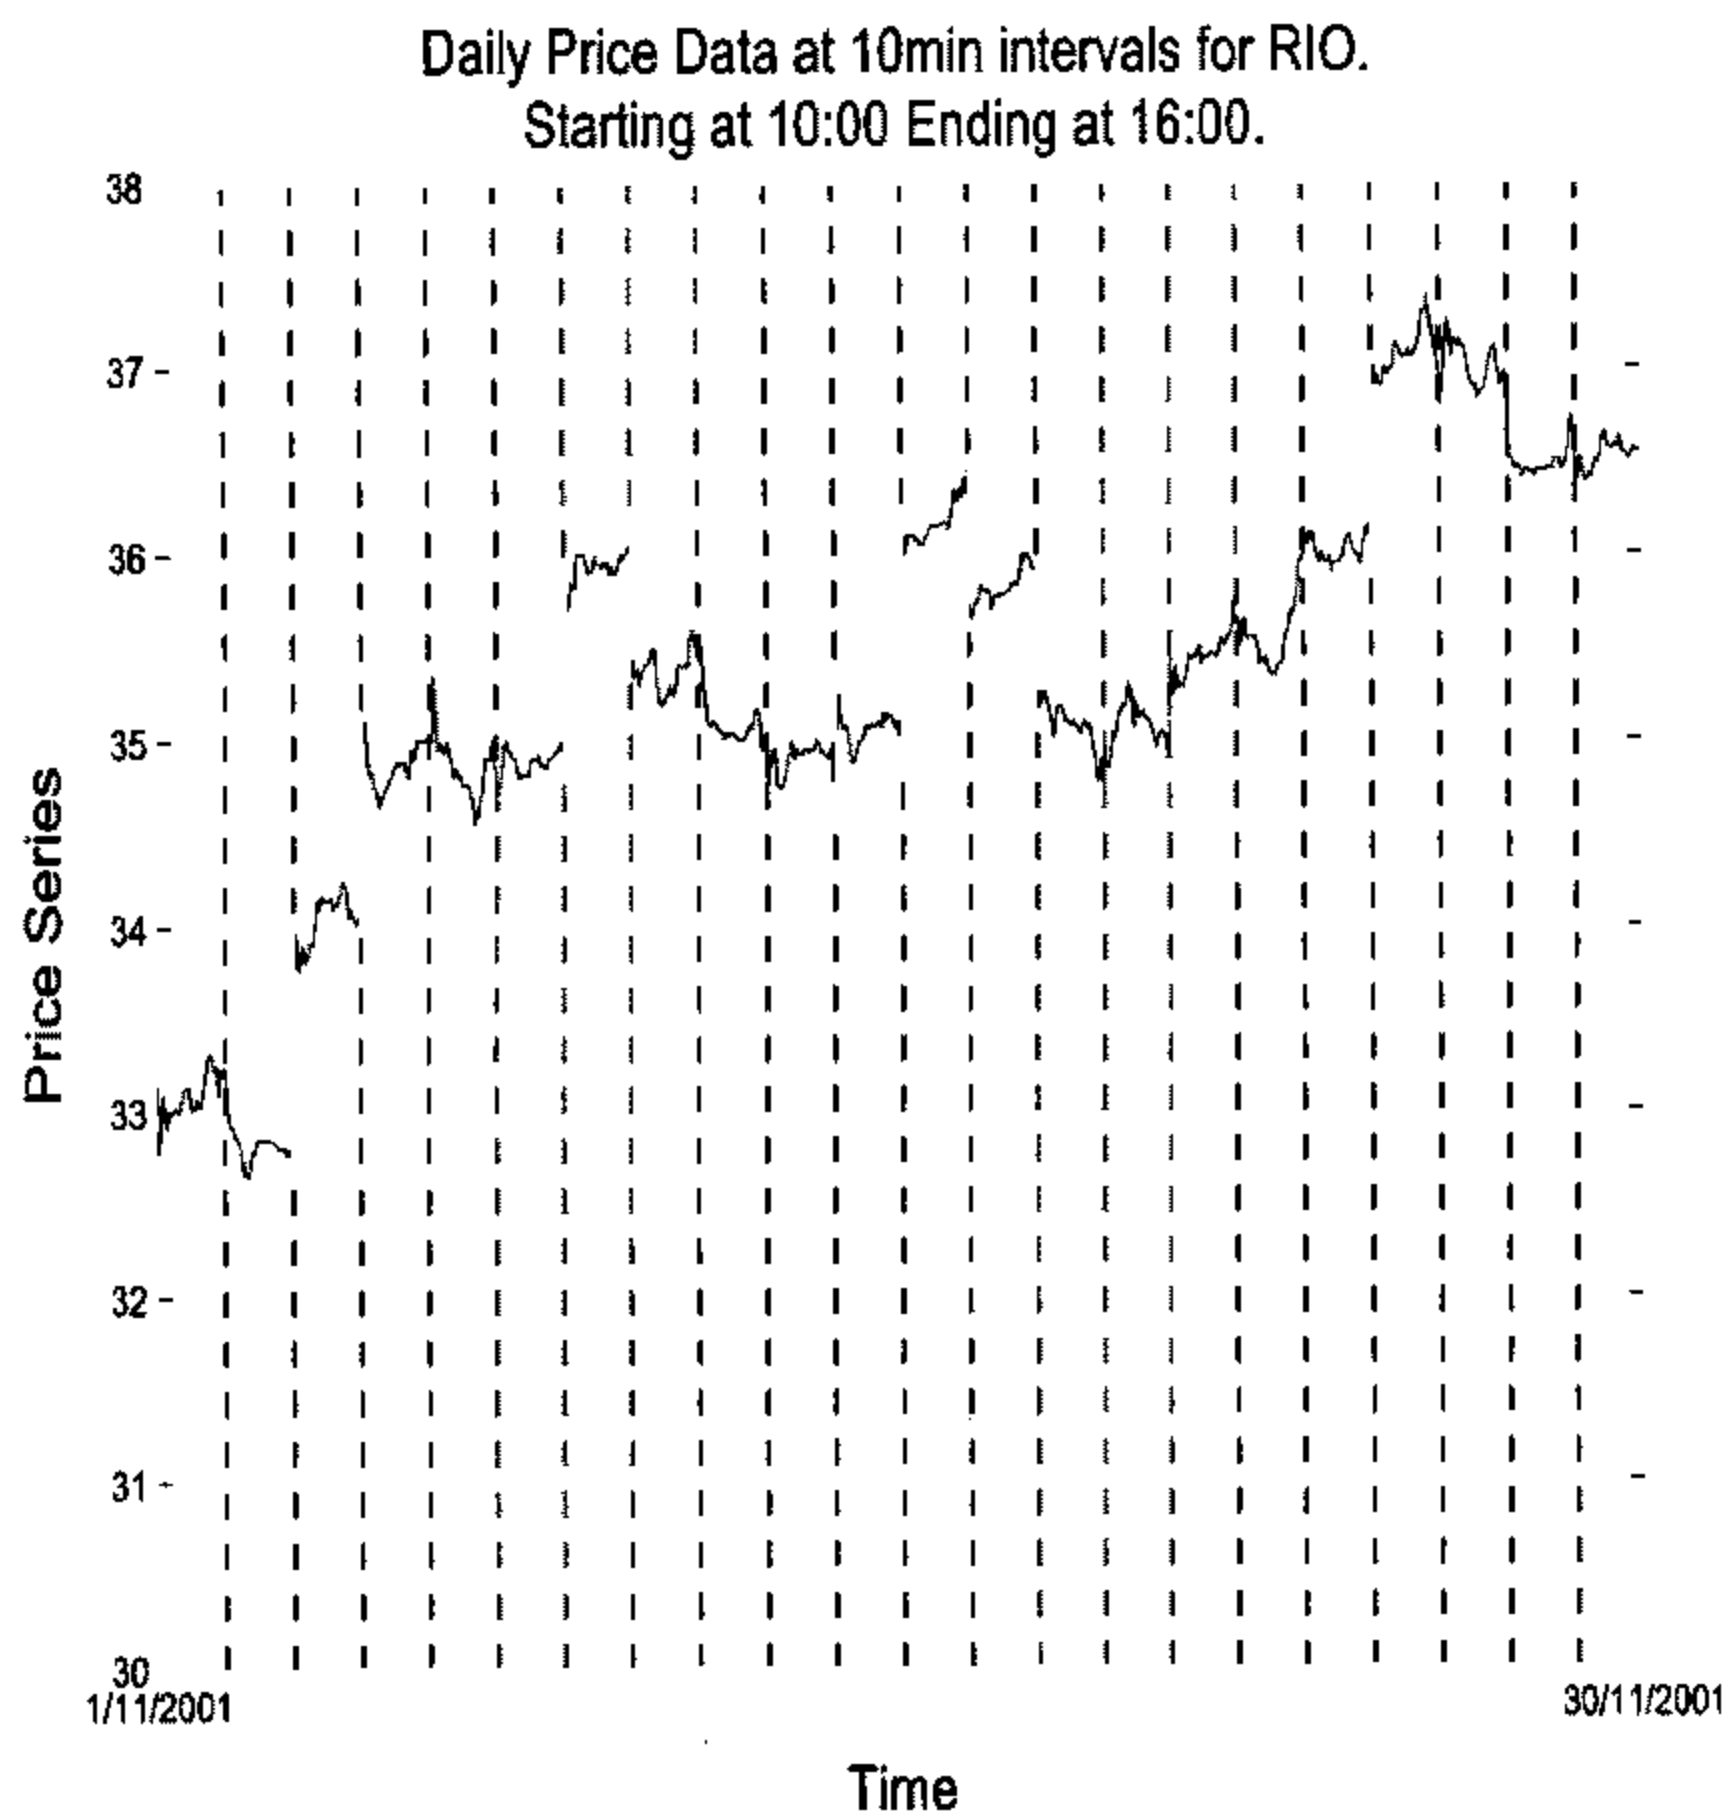


Figure 3.1: Month of trading of Rio Tinto (RIO). The dashed lines correspond to the opening and closing times. Apparent in this plot are the large jumps in price taking place between the close and open of consecutive days.

We propose that an equity return on the Australian Stock Exchange consists of two processes, one that drives the stock during trading times and another that operates during non-trading times. The implication is that because Australia is a relatively small market, the general behaviour of the ASX will be influenced mainly by the overnight values since any major market shifting information will arrive from the large markets of US and Europe while Australian markets are closed. It is apparent that the overnight process is discrete and should not be modelled by a continuous process. We term this process the 'overnight jump process'. In

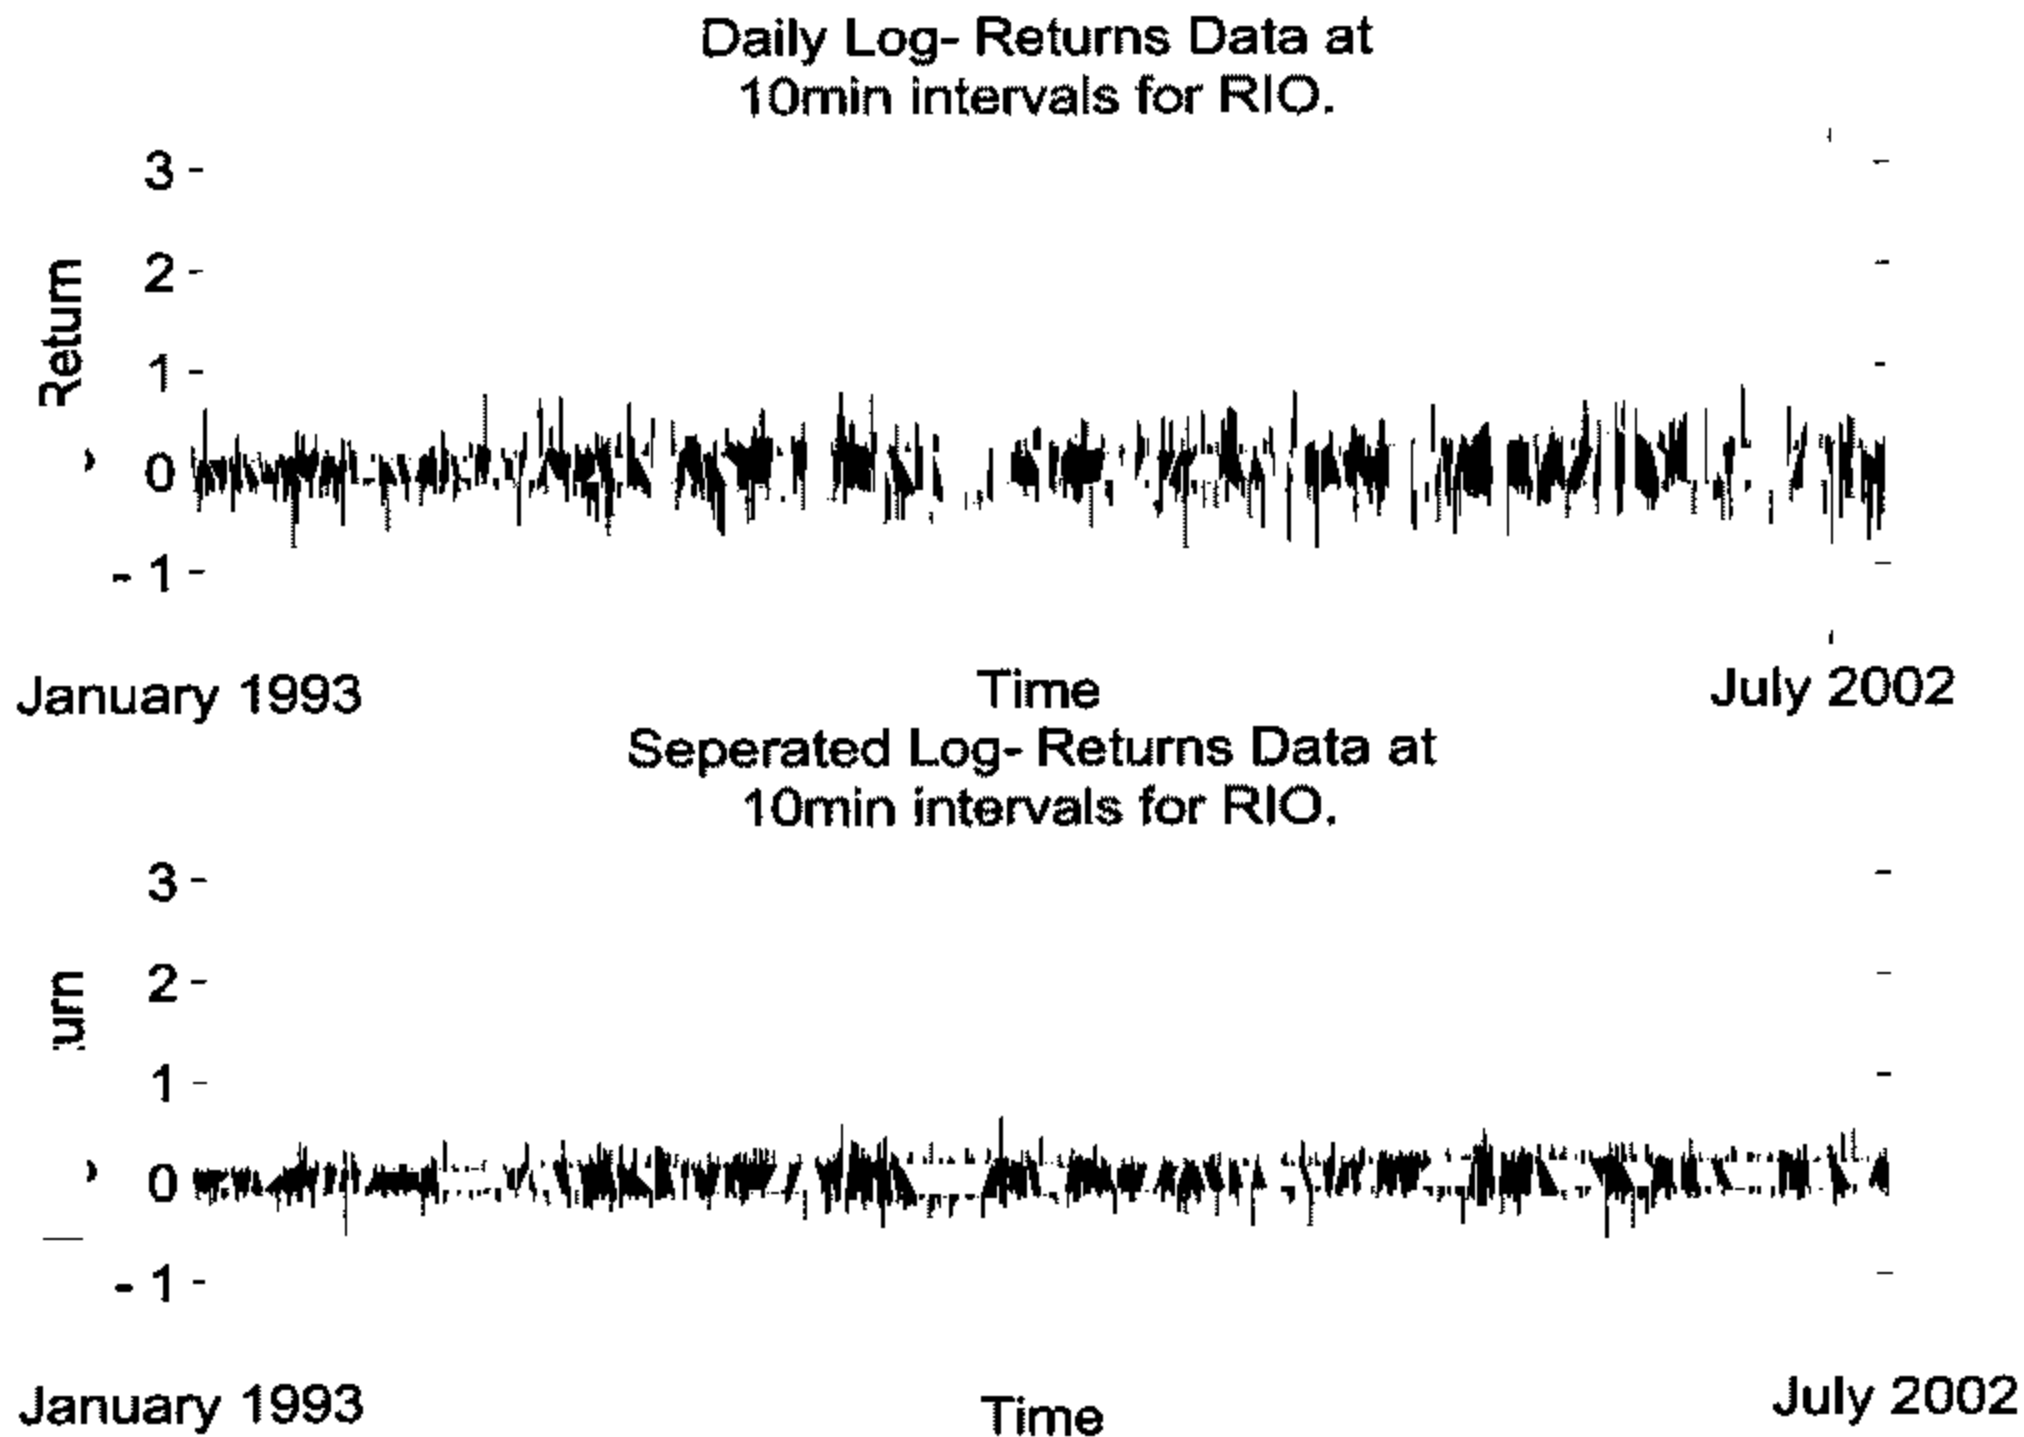


Figure 3.2: The log-return process $L_{t,\tau}$ for Rio Tinto (RIO) from 1993 to 2002 calculated in for 10 minute intervals. The top picture shows the return process including the overnight log-return, while the bottom picture has had the overnight log-return removed.

comparison to the overnight jump process, the intraday traded process is more like a continuous time process. It has a multitude of scales and looks more like a classical random noise process. Consequently we term it the ‘intraday noise process’. Hence when we look at daily price series we see the total stock price consisting of mainly the overnight jump process plus a smaller contribution due to the intraday noise.

3.2.2 Intraday trend and master volatility measure

Turning our attention to the intraday noise process, we find that this process contains some remarkably consistent behaviour. Averaging $|L_{t,\tau}|$ at each intraday interval gives us a measure of the intraday volatility over a day. We observe a well defined intraday volatility pattern emerging, as shown in Fig. 3.3. The volatility starts high

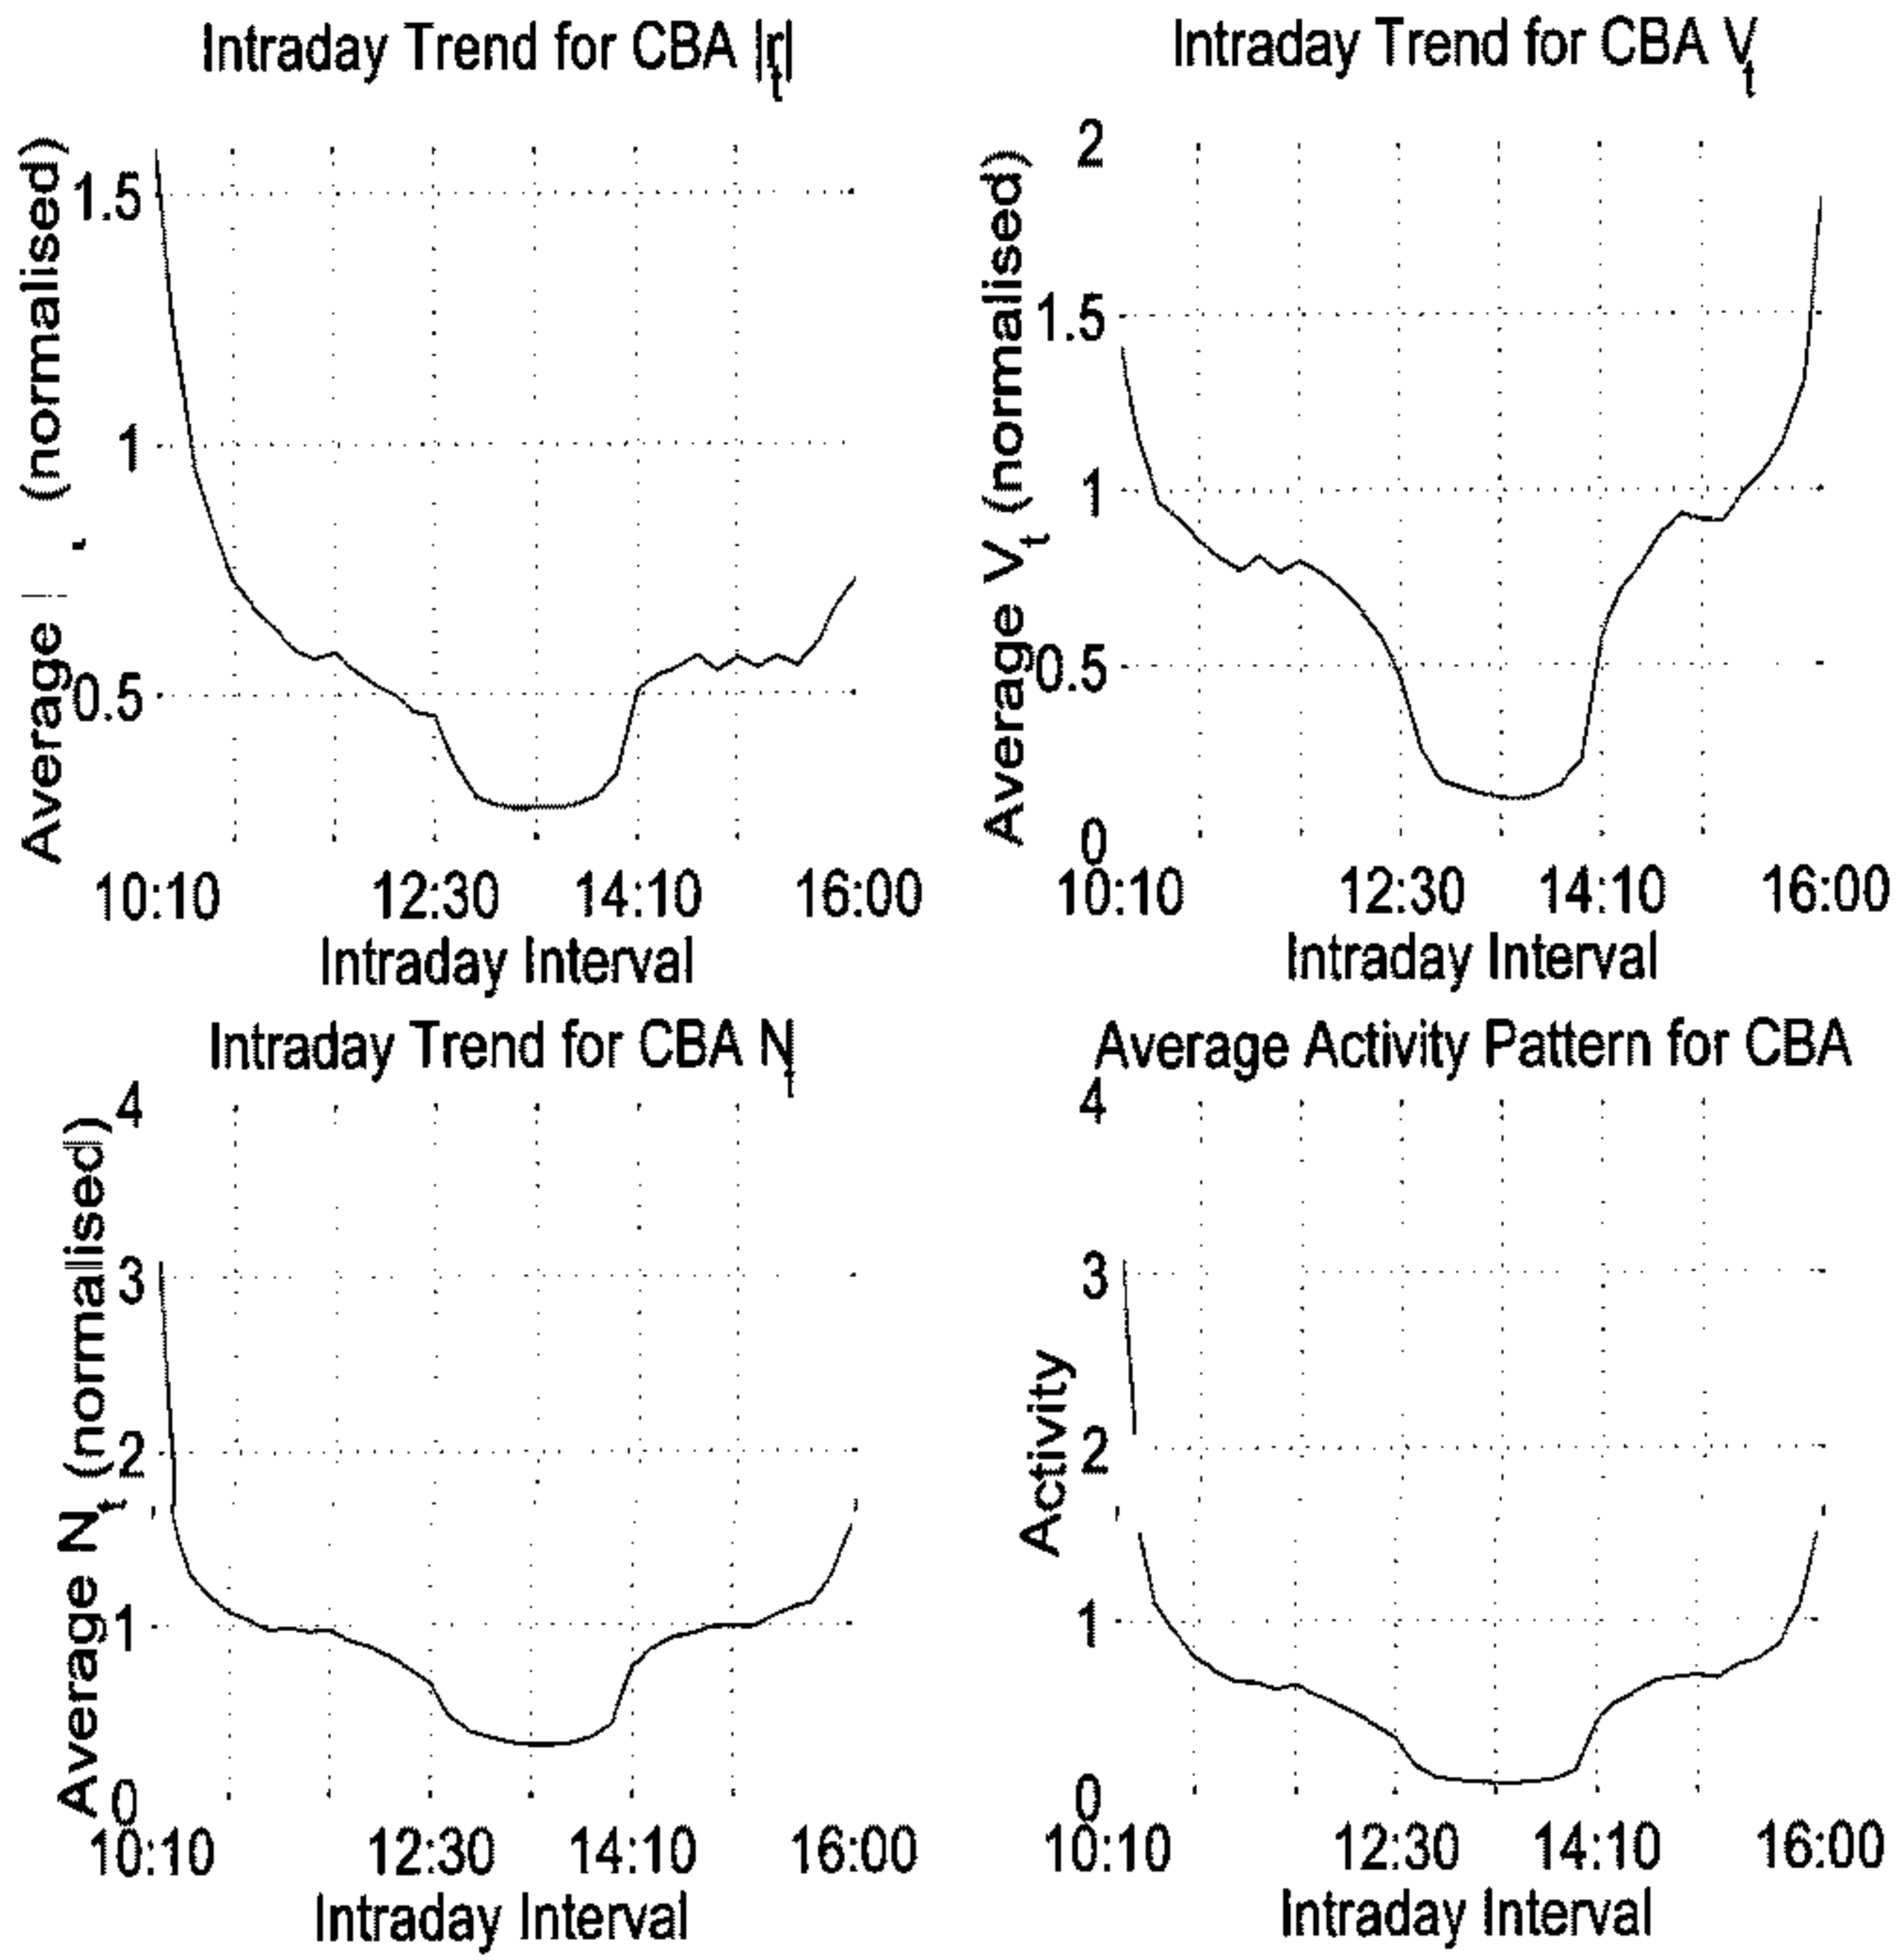


Figure 3.3: The intraday trends found in $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$, after averaging over 10 years (≈ 2500 days) of Commonwealth Bank (CBA) data. Also shown in the bottom right of the figure is the pattern of average market activity represented by the master volatility measure. This measure is a combination of the three intraday trends for $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ extracted from Principal Component Analysis.

at the opening and drops off rapidly over the morning as fund managers move quickly to correct their positions due to the overnight jump in information. Conversely, in the afternoon the volatility rapidly increases in anticipation of the close. The above pattern will also be reinforced by the presence of the many so called 'day traders' on the ASX, whose practice it is to close out all their positions at the end of each trading day and reopen their positions the following morning. The rationale of day traders is to avoid overnight exposure to risk. Interestingly this plot also provides us with a picture of the social behaviour of ASX equity traders. The volatility can be seen to drop off suddenly after the interval 12.20-12.30pm and pick up again on the interval 14.00-14.10pm. These times correspond to the close of options trading on the ASX and is typically the preferred lunchtime of most traders. Examination of the average volume traded or the average transaction frequency during each intraday interval yields similar measures for the intraday trading activity, shown in Fig. 3.3. Similar behaviour has been found in studies conducted on both the New York and London Stock Exchanges by Benston and Hagerman (1974), Wood, McInish and Ord (1985) and Abhyankar et.al. (1997).

The time series $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ are of much interest to financial researchers as they are proxies for the market volatility and hence can be thought of as representing the amount of information, or activity, in the market. Granger and Ding (1996), Clark (1973), Mandelbrot and Taylor (1967), and Geman and Ane (1997) have proposed these series as measures for volatility. The three intraday trends of Fig. 3.3 can be interpreted as measuring the average level of market activity over the course of a day. While each pattern takes on roughly the same shape, each also contains different information concerning the state of the market. It would be of great benefit to be able to summarise the results of each intraday pattern into a single measure for market activity.

To combine the three intraday patterns into a single measure for average market activity we use the method of principal component analysis (PCA) as described by Jolliffe (1986). The main idea of PCA is to obtain a set of m new variables that summarise most of the variation present in the set of n original variables, with $m \ll n$. Expanding on previously detected scaling behaviour of financial markets in

Bertram (2004), we take the principal components to be of the form,

$$\mathbf{z}_k = x_1^{\alpha_{1,k}} x_2^{\alpha_{2,k}} x_3^{\alpha_{3,k}} \quad (3.2)$$

where $\alpha_{1,k}, \alpha_{2,k}, \alpha_{3,k}$ are the k^{th} PC coefficients and x_1, x_2, x_3 are the volatility functions for $|L_{t,\tau}|, V_{t,\tau}$ and $N_{t,\tau}$. The coefficients of, and the variation explained by, the corresponding PC's for a typical stock in the data set are presented in Table B.1. On examination it is seen that the first PC accounts for 95.83% of the total variation of the data while the second and third PC's collectively account for less than 5%. We term the first principal component the 'master volatility measure' (MVM), and it represents the best 1-dimensional summary of the information contained in the three intraday patterns. This new measure gives the average market activity trend seen in Fig. 3.3.

The presence of this strong periodic component means that much of the interesting behaviour will be swamped by these trends. For example, Fig. 3.4 shows the autocorrelations of the absolute return, volume and transaction frequency. We see in this figure the extent to which the periodic trend prevents us from seeing any other behaviour that may exist in the autocorrelation function. Also note the spectral density in Fig. 3.5 which is again dominated by periodic singularities caused by the intraday trends.

3.2.3 Business time

The strong periodic nature of intraday financial time series results in these effects masking other interesting phenomena. In order to examine the data for other effects we first take these deterministic trends into account by de-seasonalising the time series. Several researchers have investigated methods to remove and or account for deterministic trends such as those shown in Fig. 3.3, Fig. 3.4 and Fig. 3.5. A special variety of discrete time series models known as P-GARCH has even been proposed by Bollerslev and Ghysels (1996) in order to address this effect.

When dealing with the log-return, $L_{t,\tau}$, the most common method is due to Taylor and Xu (1997) who calculated the variance at each intraday interval by averaging

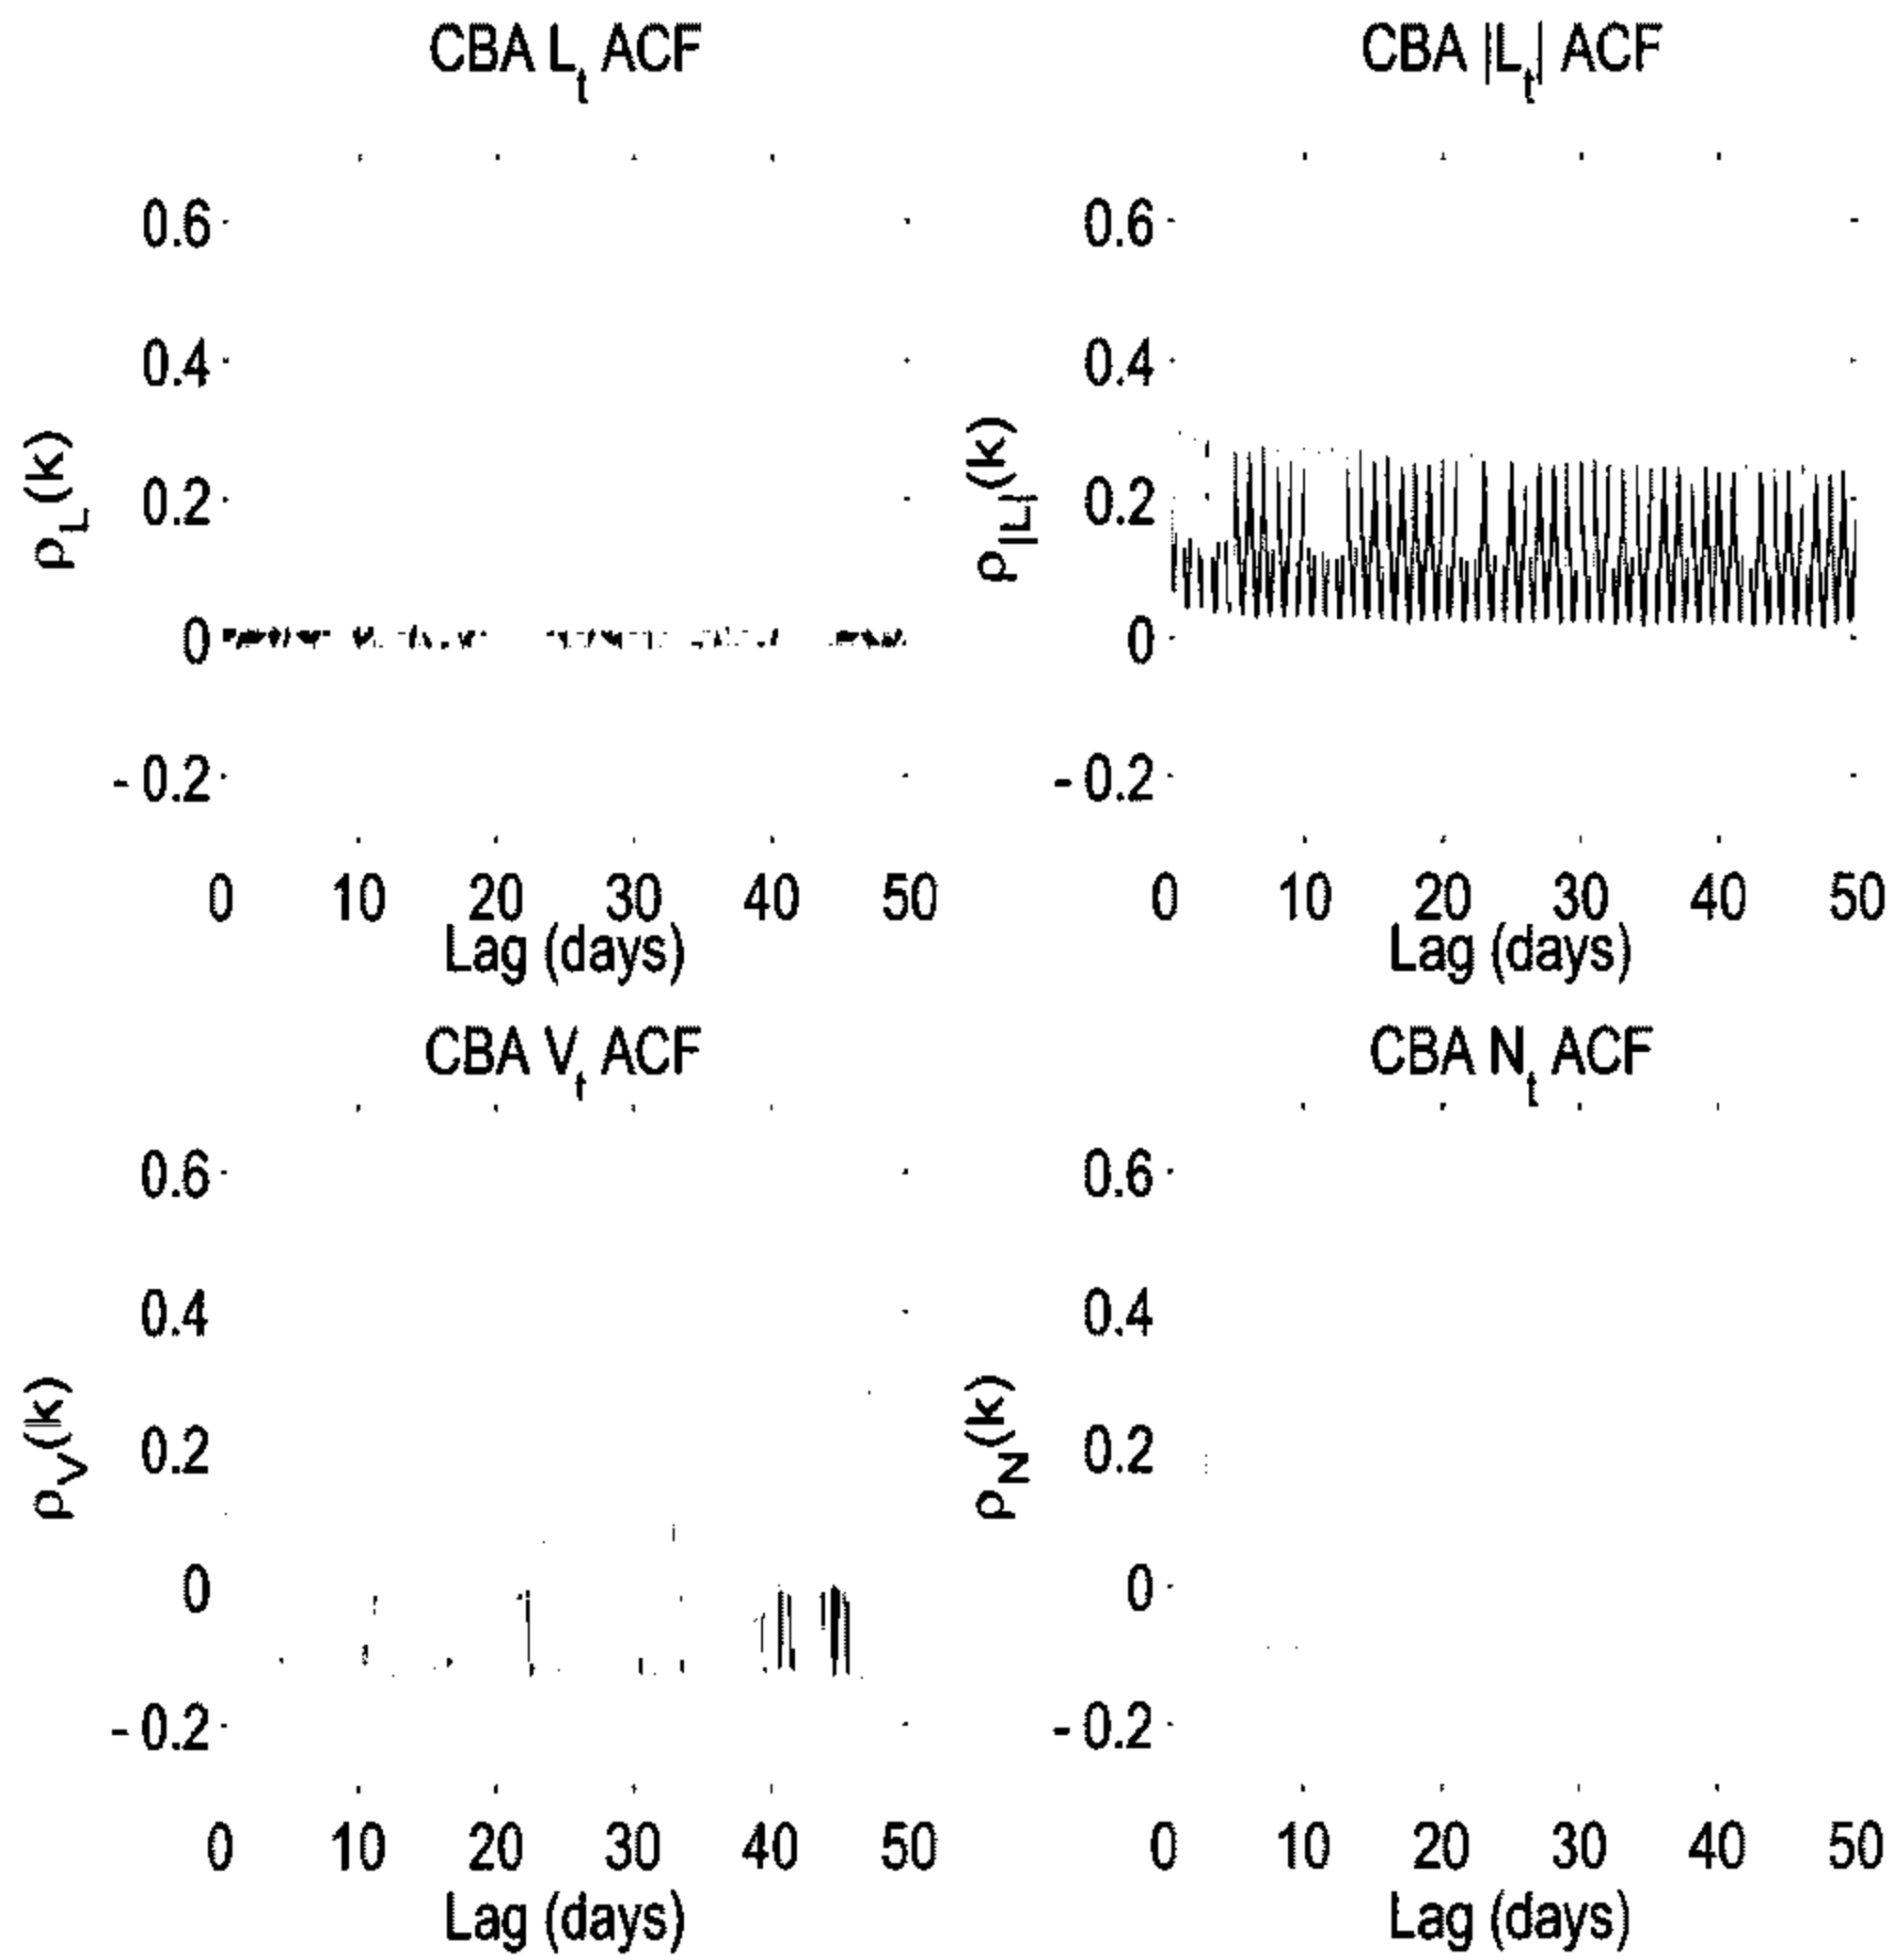


Figure 3.4: Autocorrelation of $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. We see here the strong periodic nature of the data.

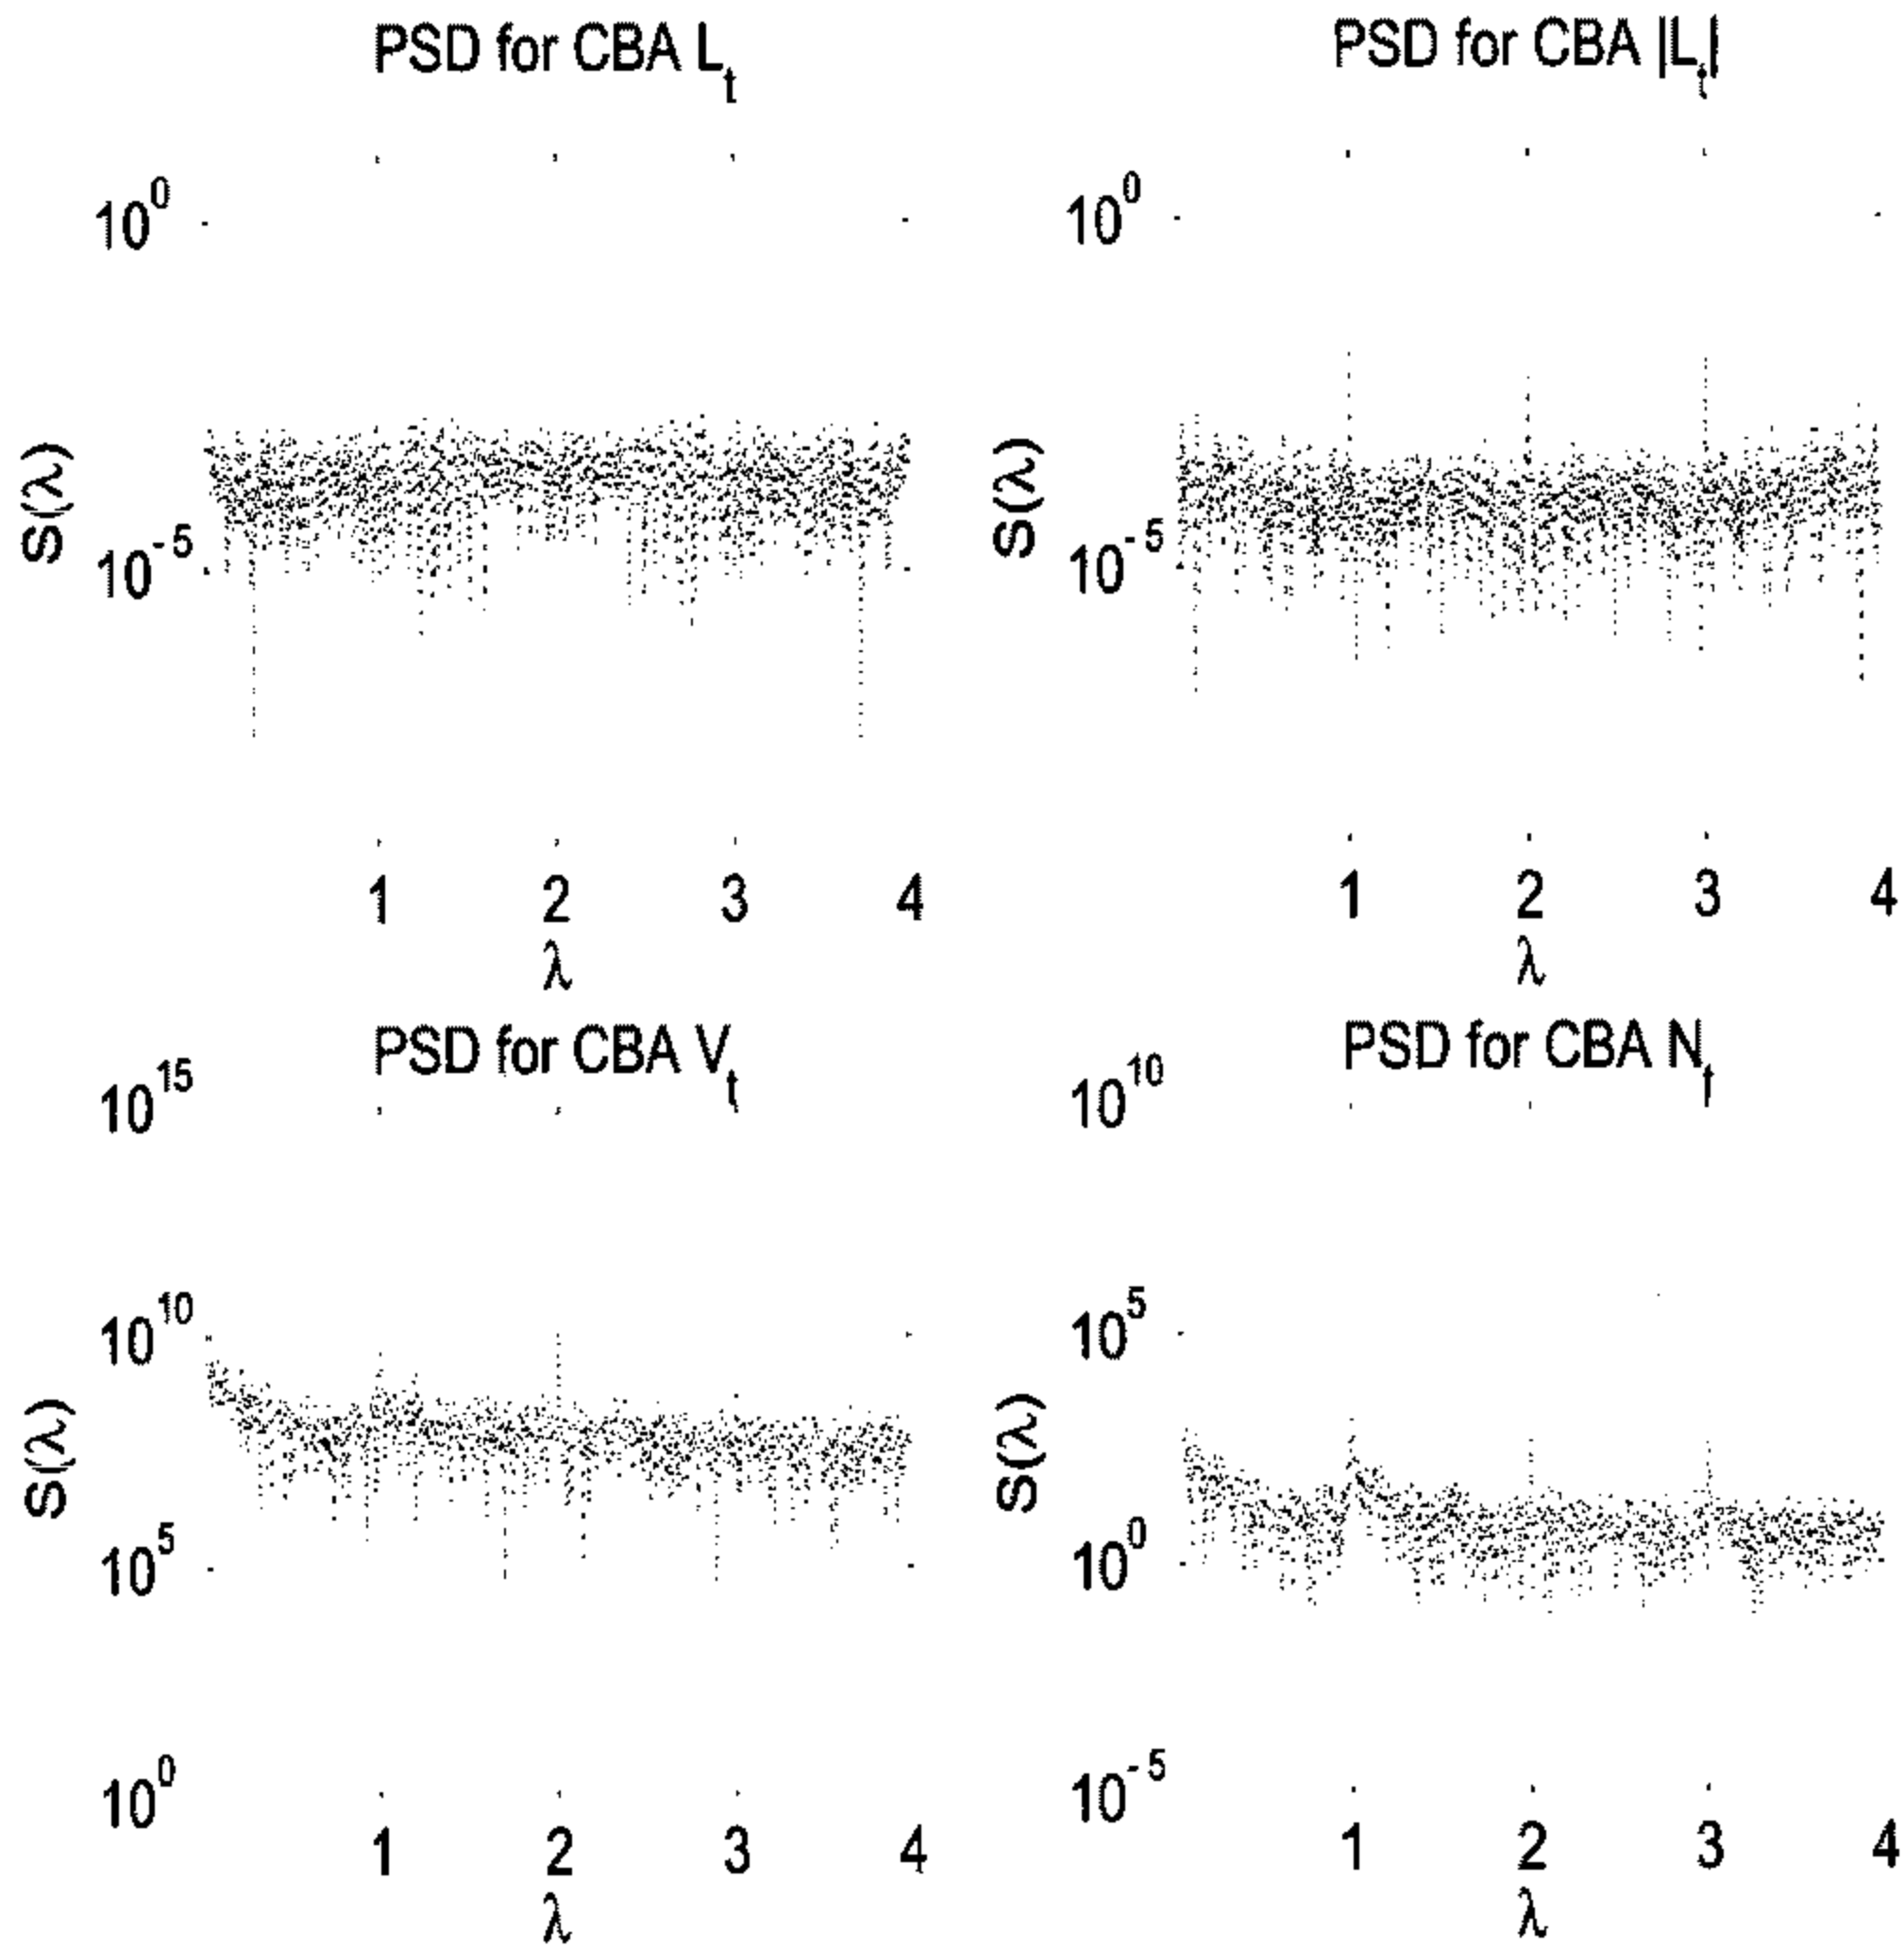


Figure 3.5: Spectral Density of $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. The strong periodic nature of the data is indicated by the spikes at daily periods.

the squared log-return across all days in the data set,

$$\sigma_i^2 = \sum_{d=1}^D L_{d,i}^2 \quad (3.3)$$

with $i = 1 \dots N$, where N is the number of sampling intervals per day and D is the total number of days in the data set. The de-seasonalised return is defined as,

$$\tilde{L}_{t,\tau} = \tilde{L}_{i,t} = \frac{L_{i,t}}{\sigma_i}. \quad (3.4)$$

Another method similar to this involves estimating the intraday volatility using the so-called flexible Fourier form (FFF) method of Anderson and Bollerslev (1997) and proceeding via Eq. (3.3). Martens, Chang and Taylor (2002) provided a comparison between the two estimation procedures described above and P-GARCH and showed their abilities to deal with intraday seasonality. However, we find that these methods are unable to remove the periodic trends found in the Australian equity data. We find that data filtered by the method of Taylor and Xu (1997) still exhibits behaviour of a considerable periodic nature.

We turn to the method described in Dacorogna et.al. (2001) to address the periodic nature of the data. This method prescribes the construction of a new time scale, known as ‘business’ or ‘market’ time, on which traders observe a non-seasonal volatility. In physical time we witness that the amount of activity taking place in the market over an interval Δt will change deterministically as the day progresses. Thus, the new time scale, $B(t)$, is defined such that the level of market activity is kept constant in each interval $\Delta B(t)$.

To construct the business time scale we can represent the activity by a rescaled volatility measure $v(t)$. Fig. 3.6 and Fig. 3.7 display the average values of the 3 volatility measures $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ evaluated for different length time intervals. Using a scaling relation of the form,

$$\mathbb{E}[v(t)] = c\Delta t^{1/E} \quad (3.5)$$

we can relate the long time average of the volatility to the size of the sampling interval. This relation is further explored in Section 3.5 where we look at the fractal

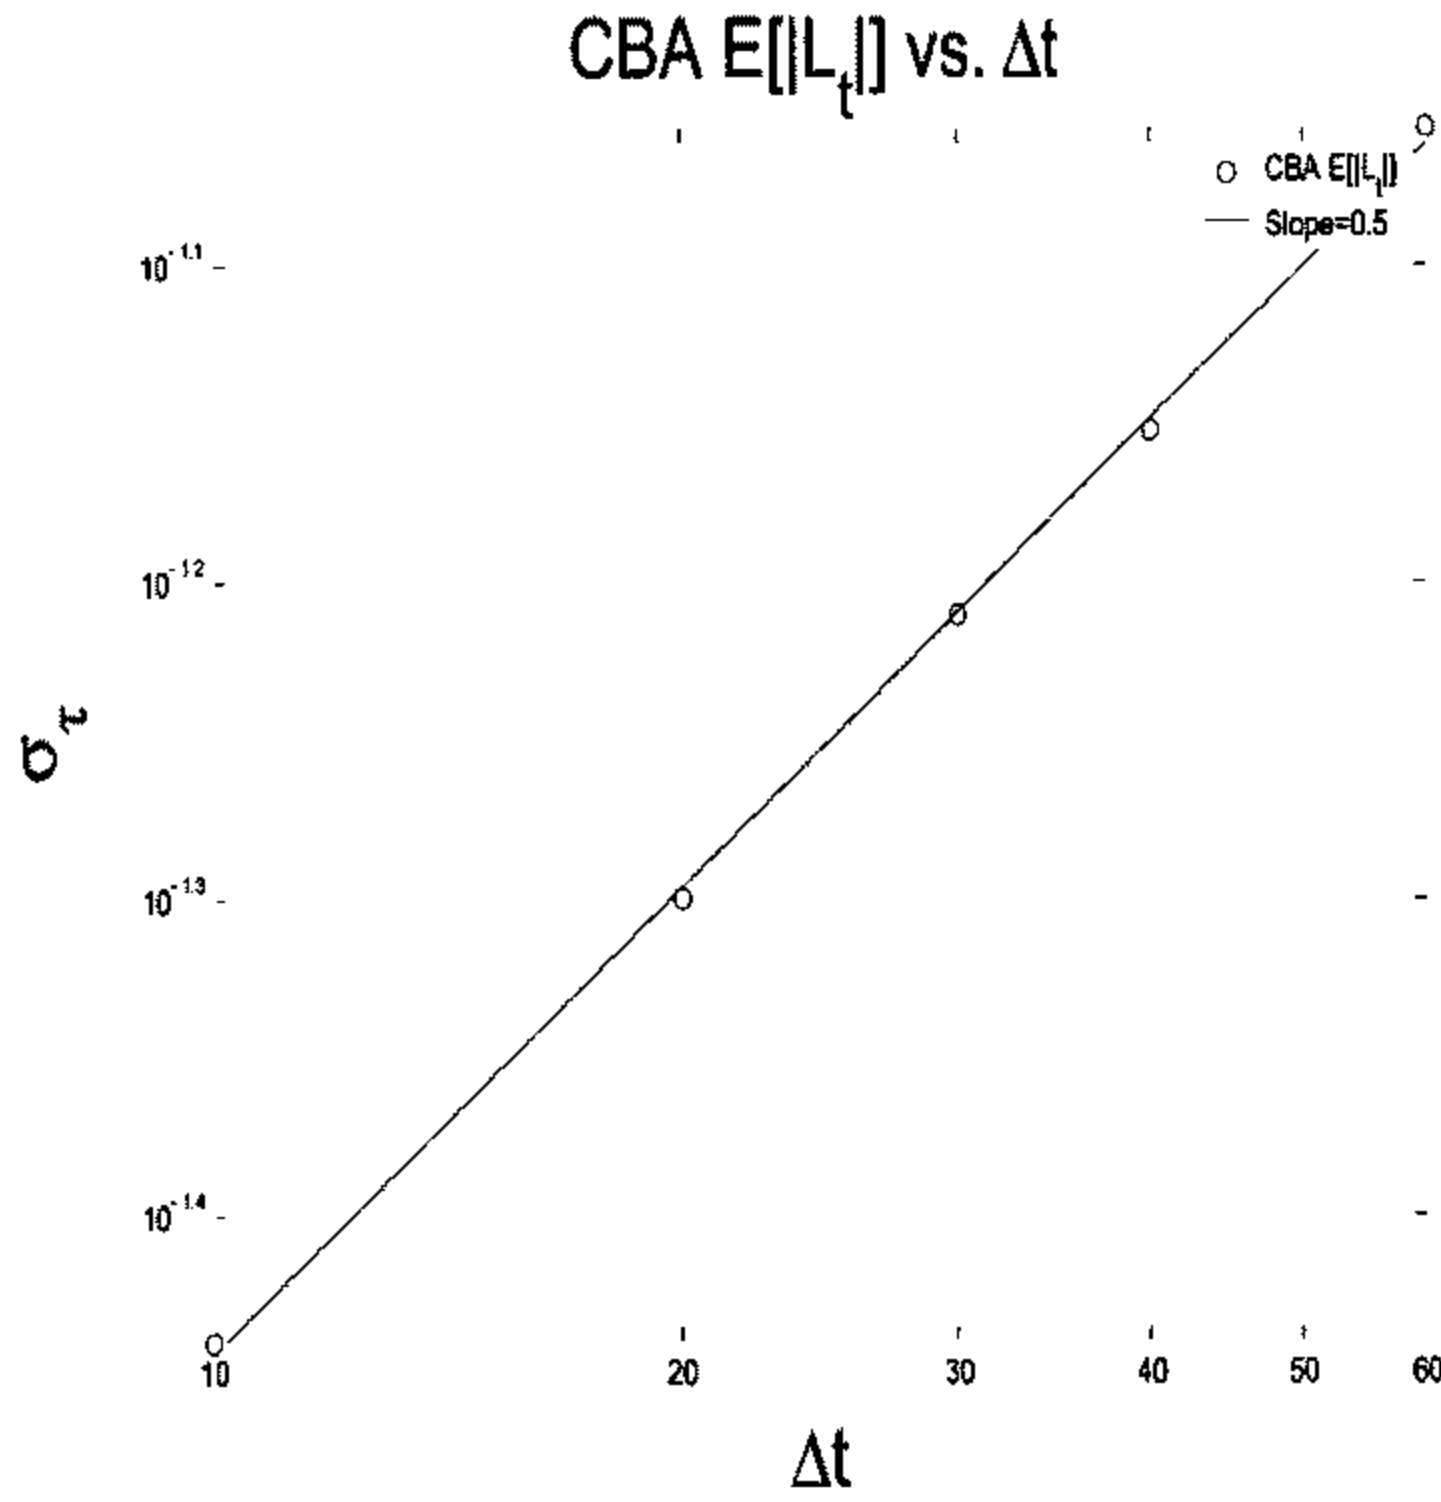


Figure 3.6: Scaled absolute returns.

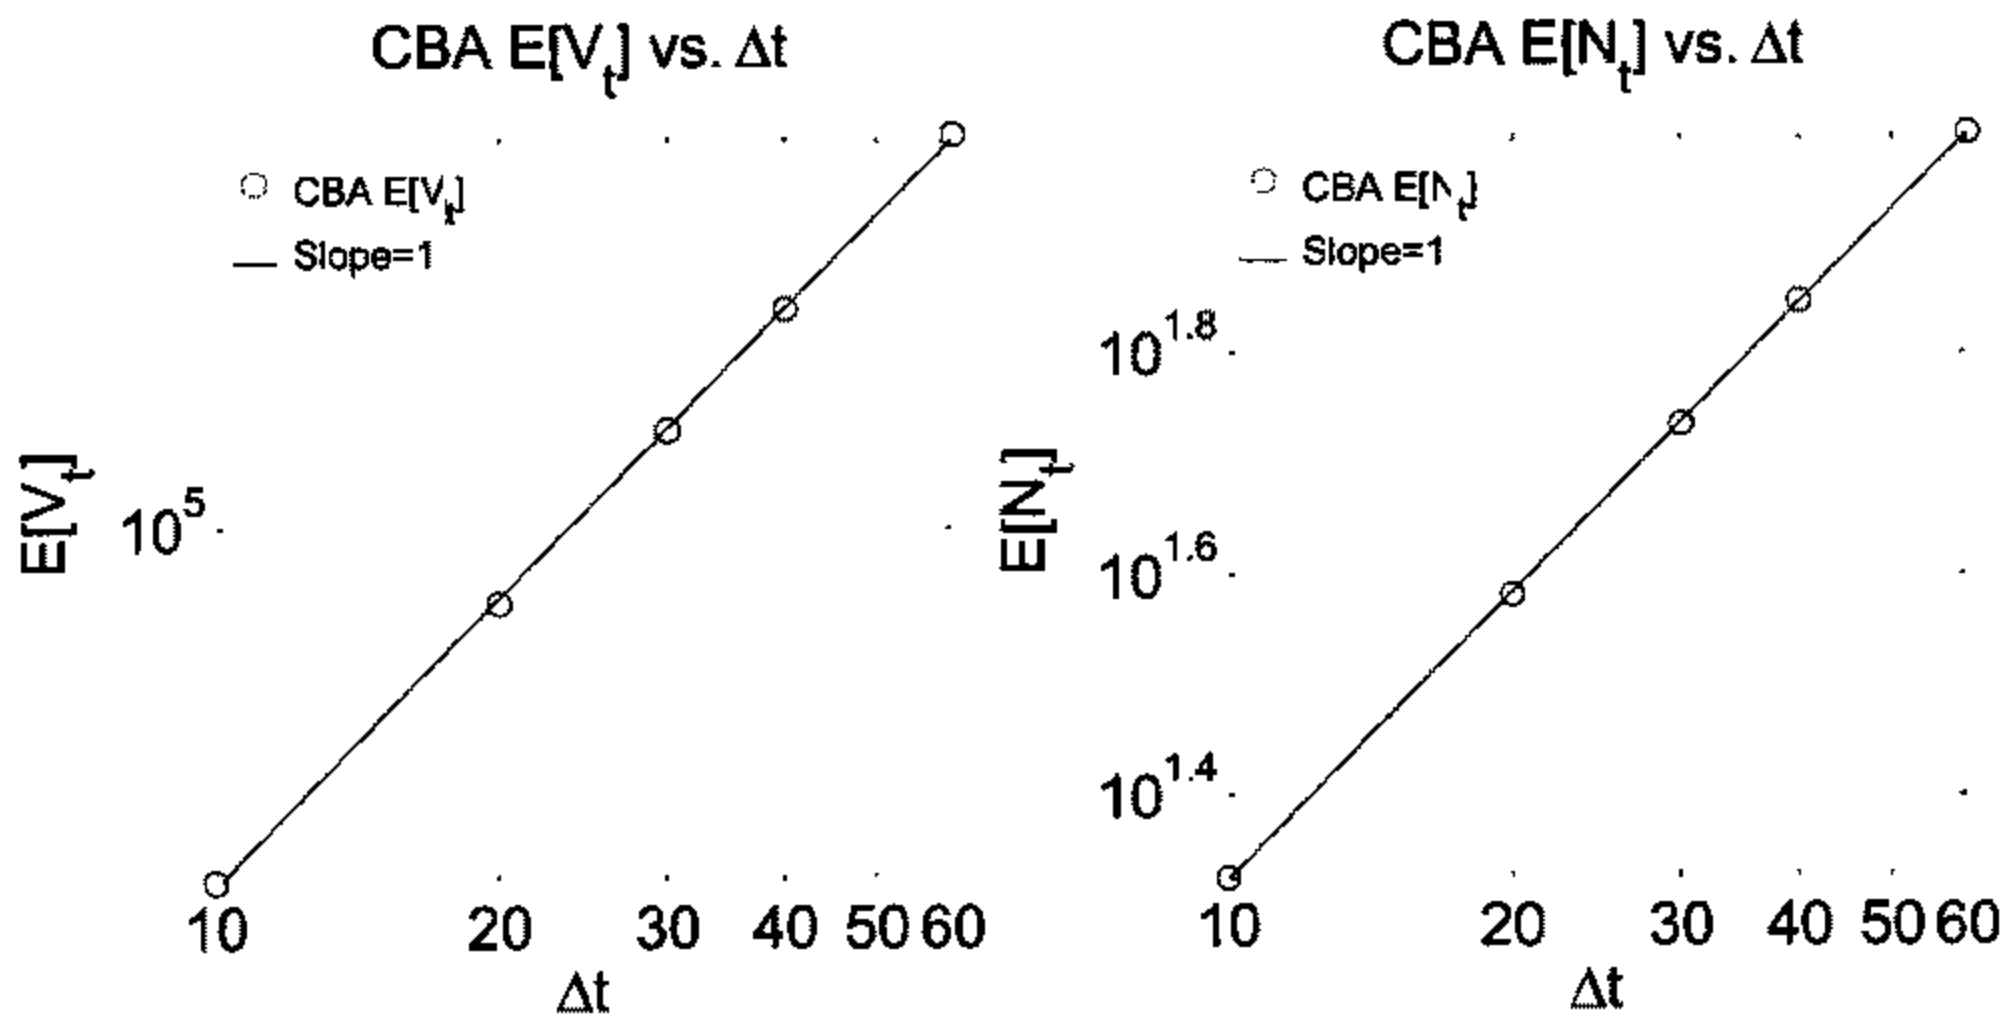


Figure 3.7: Scaled volume, $V_{t,\tau}$ and transaction frequency, $N_{t,\tau}$

properties of the data. The activity can be specified as the rescaled value,

$$a_i = \frac{1}{\Delta t} \left(\frac{\mathbb{E}[v(t_i)]}{c^*} \right)^E \quad (3.6)$$

where c^* is a constant such that

$$\Delta t \sum_{i=1}^N a_i = 1. \quad (3.7)$$

Once the activity has been constructed we then define business time as

$$B(t) = \int_0^t a(s) ds. \quad (3.8)$$

In this new time scale we choose sampling intervals $\Delta\theta(t) = \theta(t_i) - \theta(t_j)$ and invert the time scale to find the times t_i and t_j . The data can then be resampled at the times t_i, t_j . Thus the resulting sampling interval lengths will not be constant in physical time but constant in business time.

Dacorogna et.al. (2001) used this method to study foreign exchange data. They studied the scaling behaviour of the volatility in relation to Eq. (3.5) and used a polynomial to approximate Eq. (3.6). This polynomial was then integrated to obtain the required business time. Rather than using a polynomial to describe the activity, we use a cubic spline to interpolate between the sampled values of a_i . It is of interest to note that while Dacorogna et.al. (2001) found that the volatility scales as $E \approx 0.6$, our results, shown in Table B.2, illustrate that for equity data, $E \approx 0.5$ for $L_{t,\tau}$ and $E \approx 1$ for $V_{t,\tau}$ and $N_{t,\tau}$. This result highlights one of the many possible differences between foreign exchange and equity data and reinforces the need to study data of all types.

In Fig. 3.8 and Fig. 3.9 we display the business time scale constructed using $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. Here we can see the extent to which the business time scale deviates from physical time. It also shows that although the interval lengths $\Delta B(t)$ are constant in business time, their length varies in physical time as the day progresses. In the morning the time scale evolves quickly as there is much activity, so a 10 minute interval in business time may only correspond to a few seconds in physical time. Conversely, at lunch time we see that business time evolves much more slowly




Figure 3.8: Business time scale for $L_{t,\tau}$. In this figure we can see that in the morning an interval $\Delta B(t)$ corresponds to only a few minutes in physical time. However at lunchtime the same interval, $\Delta B(t)$ now corresponds over an hour of physical time.

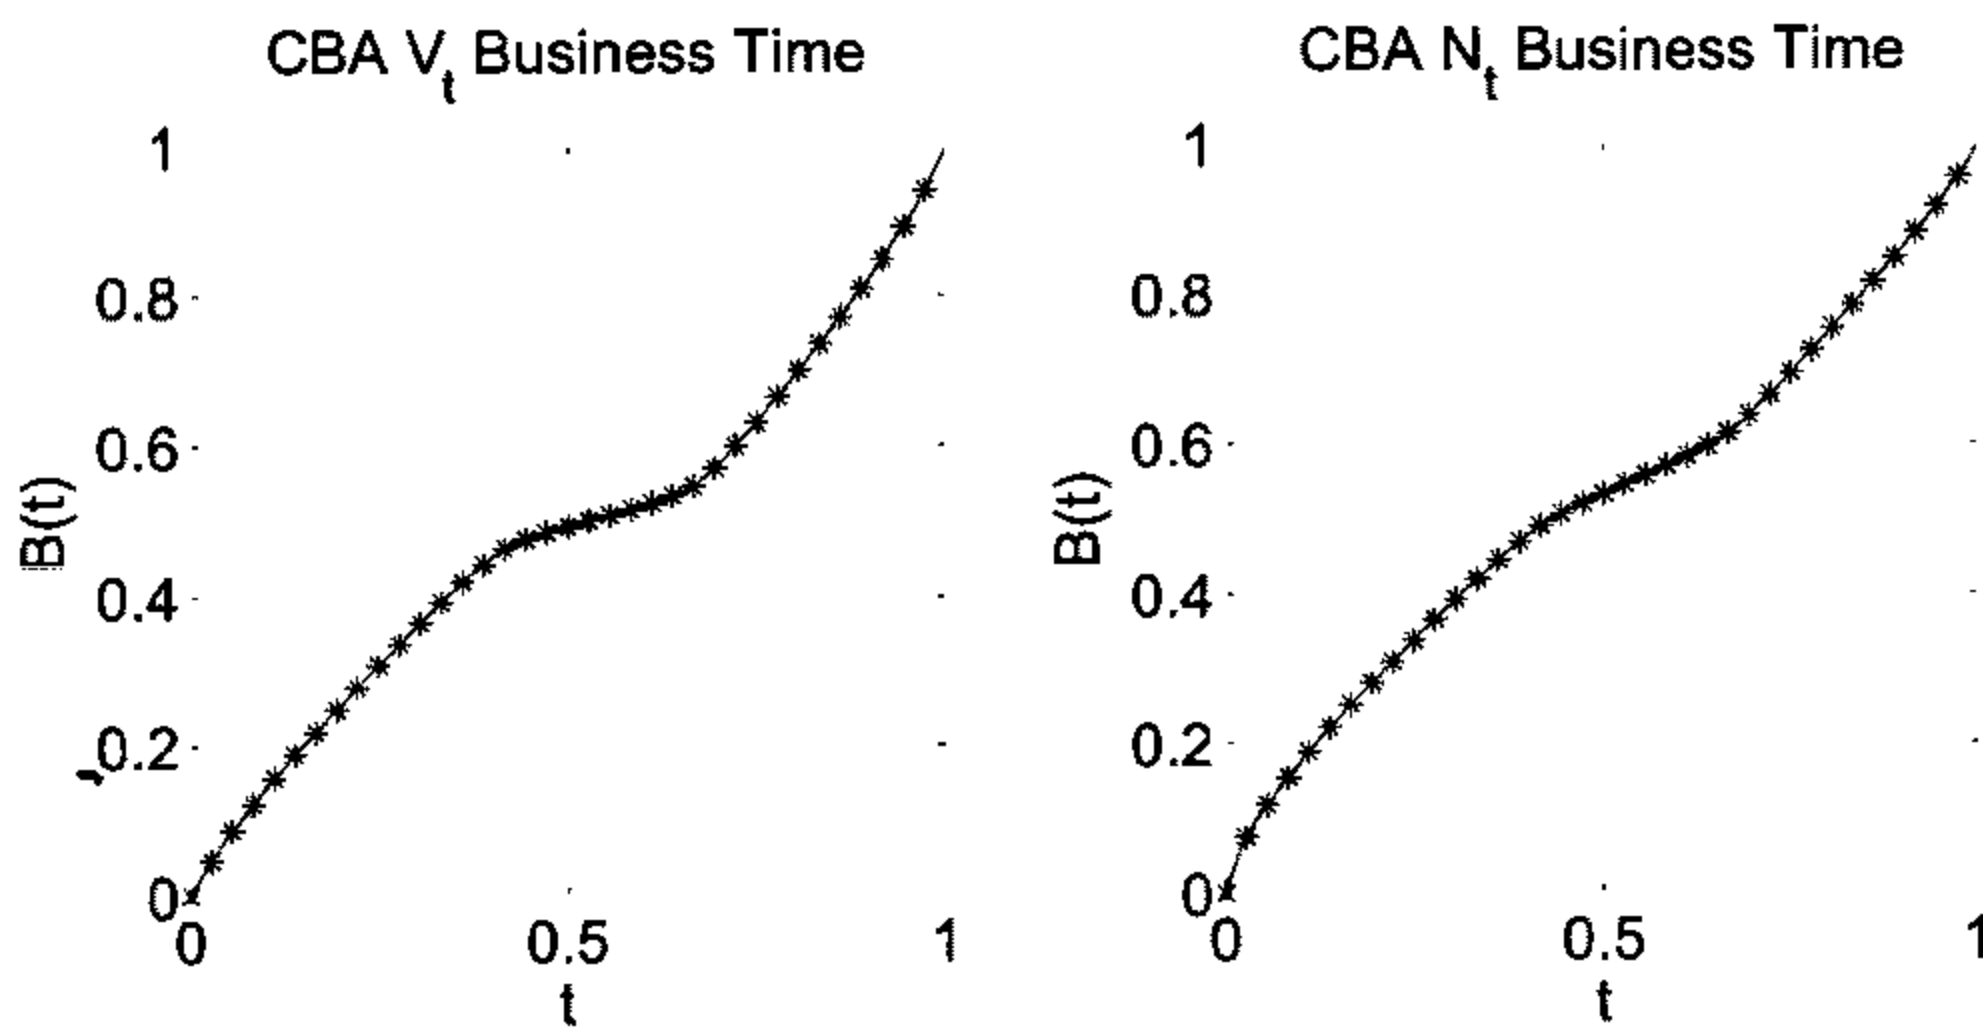


Figure 3.9: Business time scales for the volume, $V_{t,\tau}$ and transaction frequency, $N_{t,\tau}$.

than physical time and a 10 minute business time interval may correspond to around an hour of physical time.

In Fig. 3.10 we display the autocorrelation of $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ in business time. Comparing this with the previous result of Fig. 3.4 we see that the strong periodic behaviour of the time series has disappeared. Fig. 3.11 shows the spectral density of the time series in business time and it is clear in this figure that the periodic nature of the volatility has been accounted for.

3.2.4 Zero return enhancement

The distribution properties of the normalised returns have been studied extensively in recent years by many researchers. Since the work of Mandelbrot (1963) and Fama (1965) researchers in this field have been primarily concerned with the non-Gaussian nature of the distribution tails and many models have been proposed to explain this behaviour. In this section we focus our attention instead on the central region of the distribution. Examining the return distributions for stocks on the ASX we find they indicate the presence of a disproportionately large number of zero values. Fig. 3.12 shows the empirical density functions for the intraday return, calculated for a typical stock in the data set. The empirical density function displays an extremely high number of points lying in the central bin. This is the result of the log-return series containing an unusually high proportion of zero values, $\mathbf{P}[L_{t,\tau} = 0]$. This interesting property of the ASX data has also been found in high frequency DAX returns by Gorski et.al. (2002) where the authors of that paper coined the name 'zero return enhancement' (ZRE). Contrary to the findings of Gorski et.al. (2002) which state that this is due to the choice of too small a sampling interval τ , we find that for ASX data, zero return enhancement exists for quite large values of τ . Zero returns can come about in three ways: type 1. no trade taking place in the sampling interval, type 2. trading at a constant price across the interval and type 3. starting and ending the interval at the same price with a price deviation in between. Owing to the comprehensive nature of the data set it is possible to investigate the proportion of the log-return series consisting of each type of zero return. Analysis performed on different stocks shows that for $\tau = 10$ minutes, zero returns of type

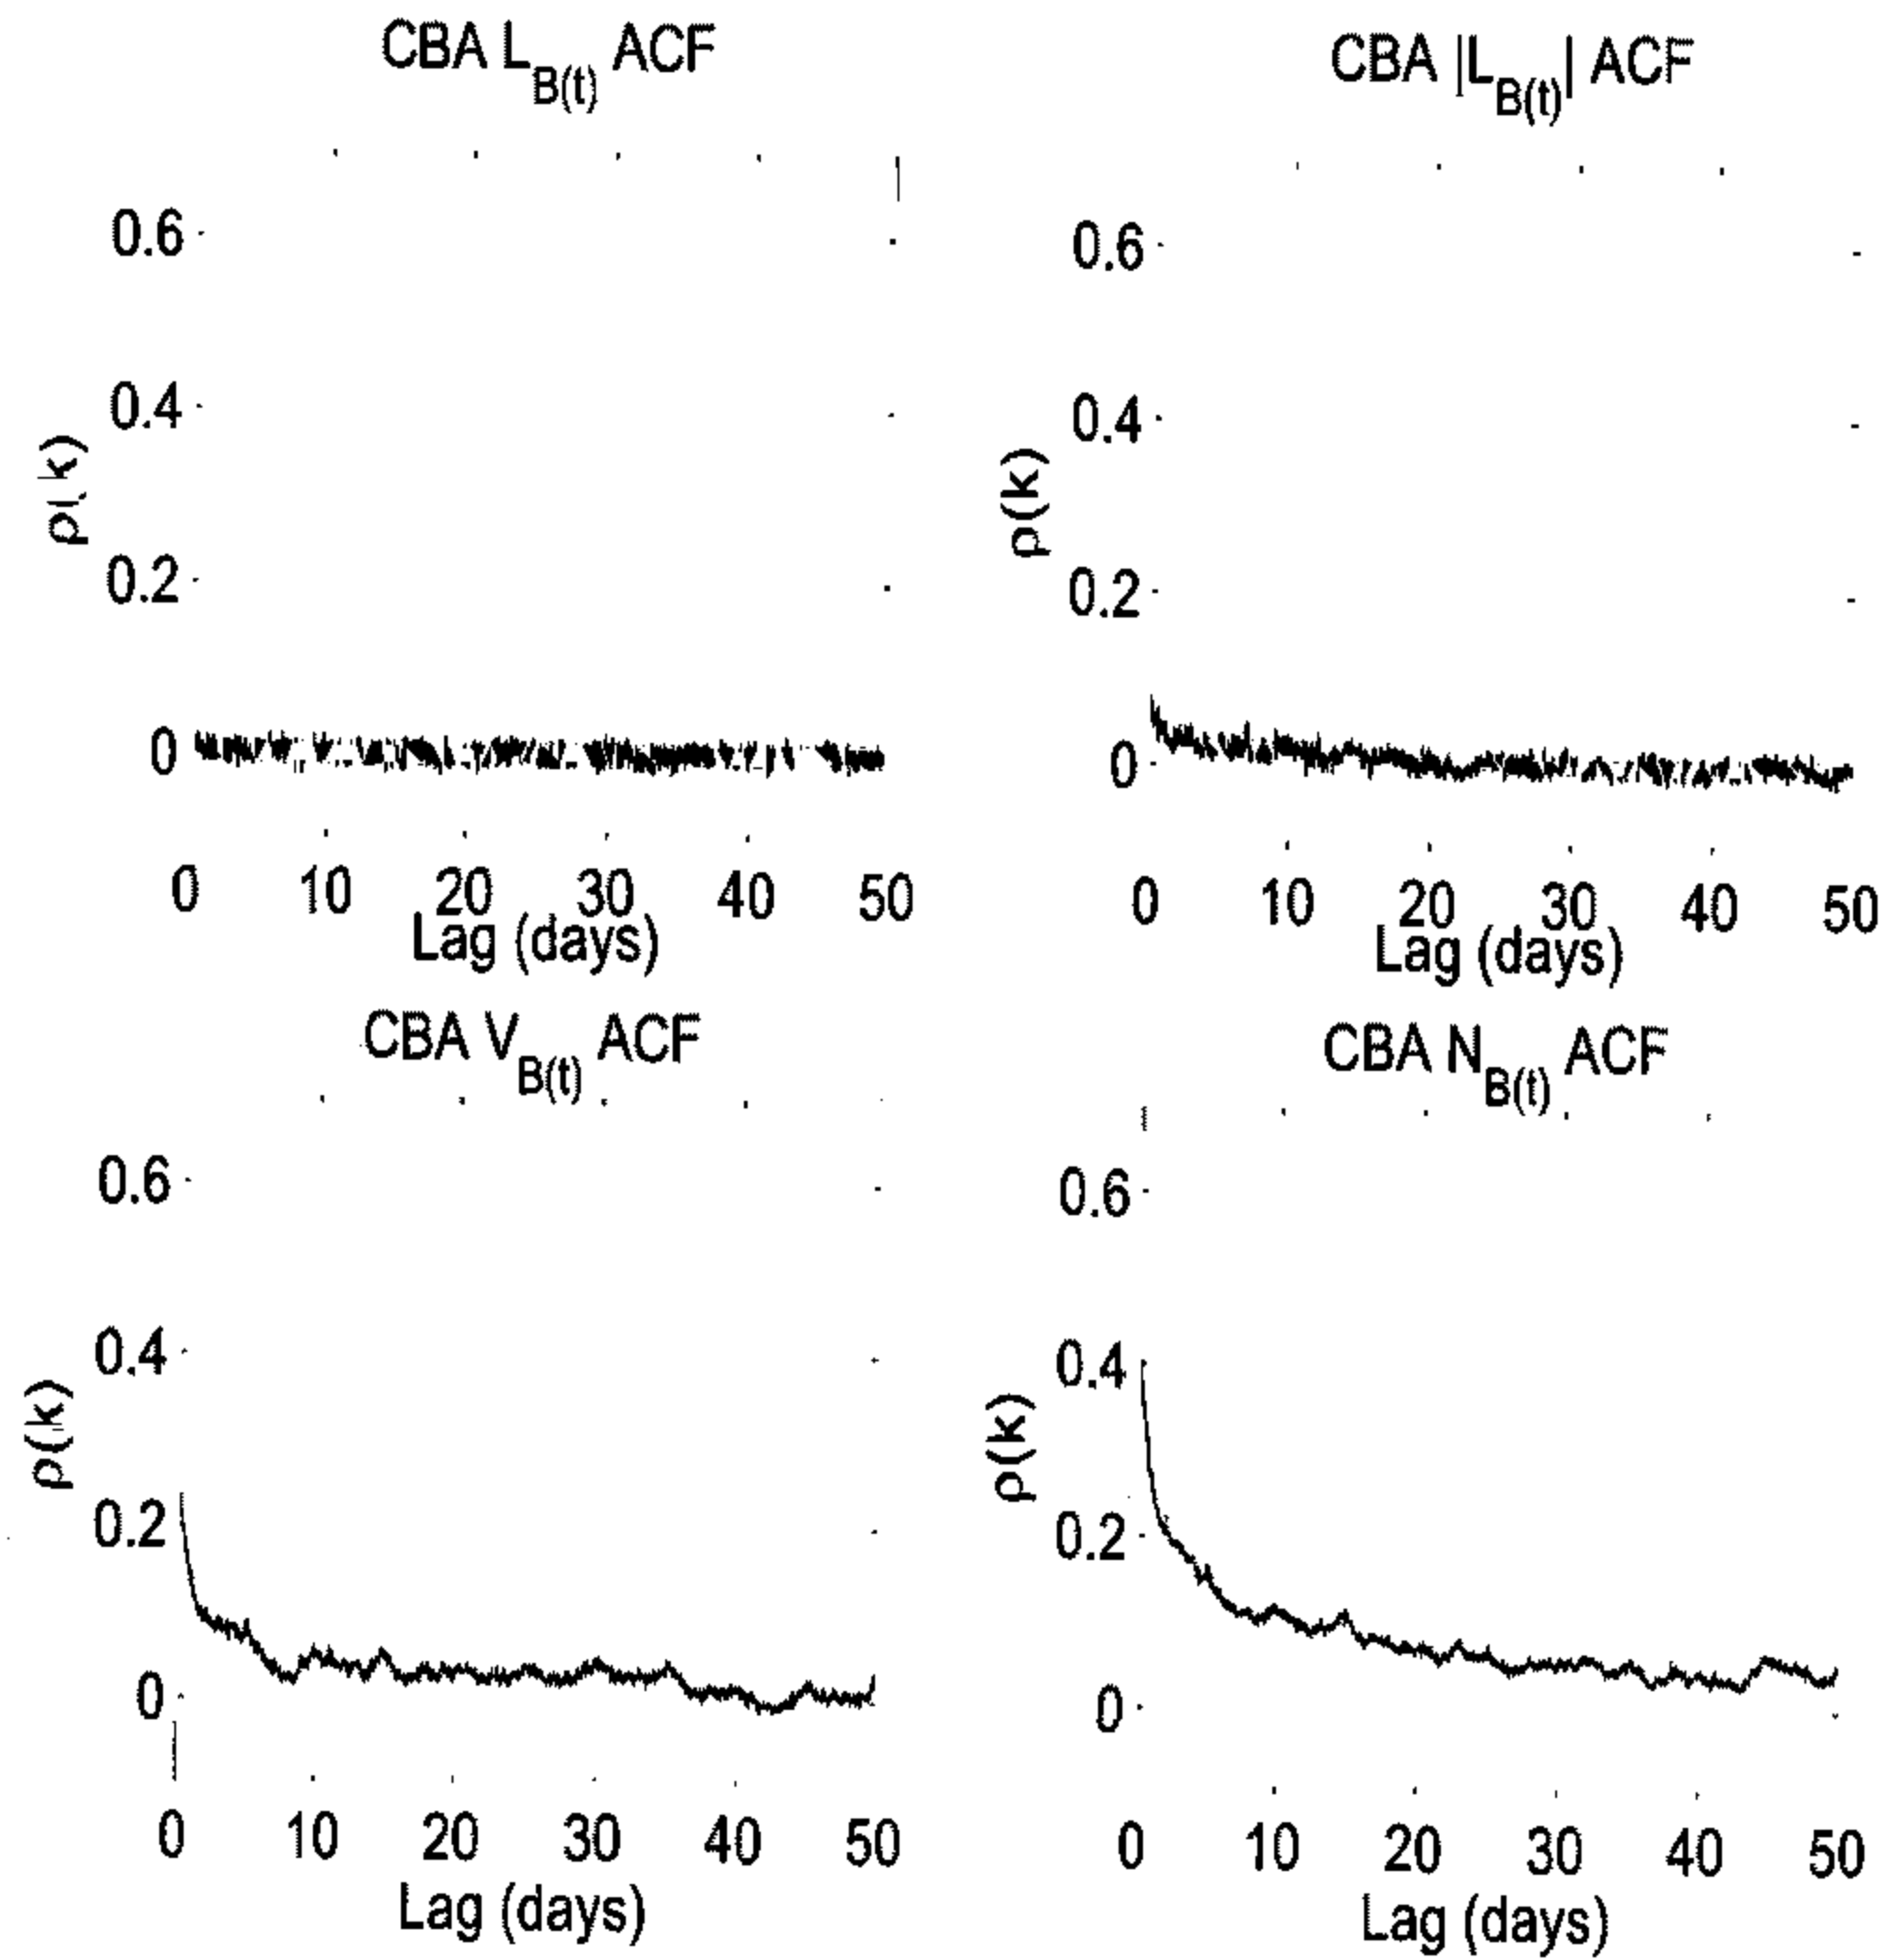


Figure 3.10: Autocorrelation of $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ in business time. We can see that the periodic nature of the data has been accounted for.

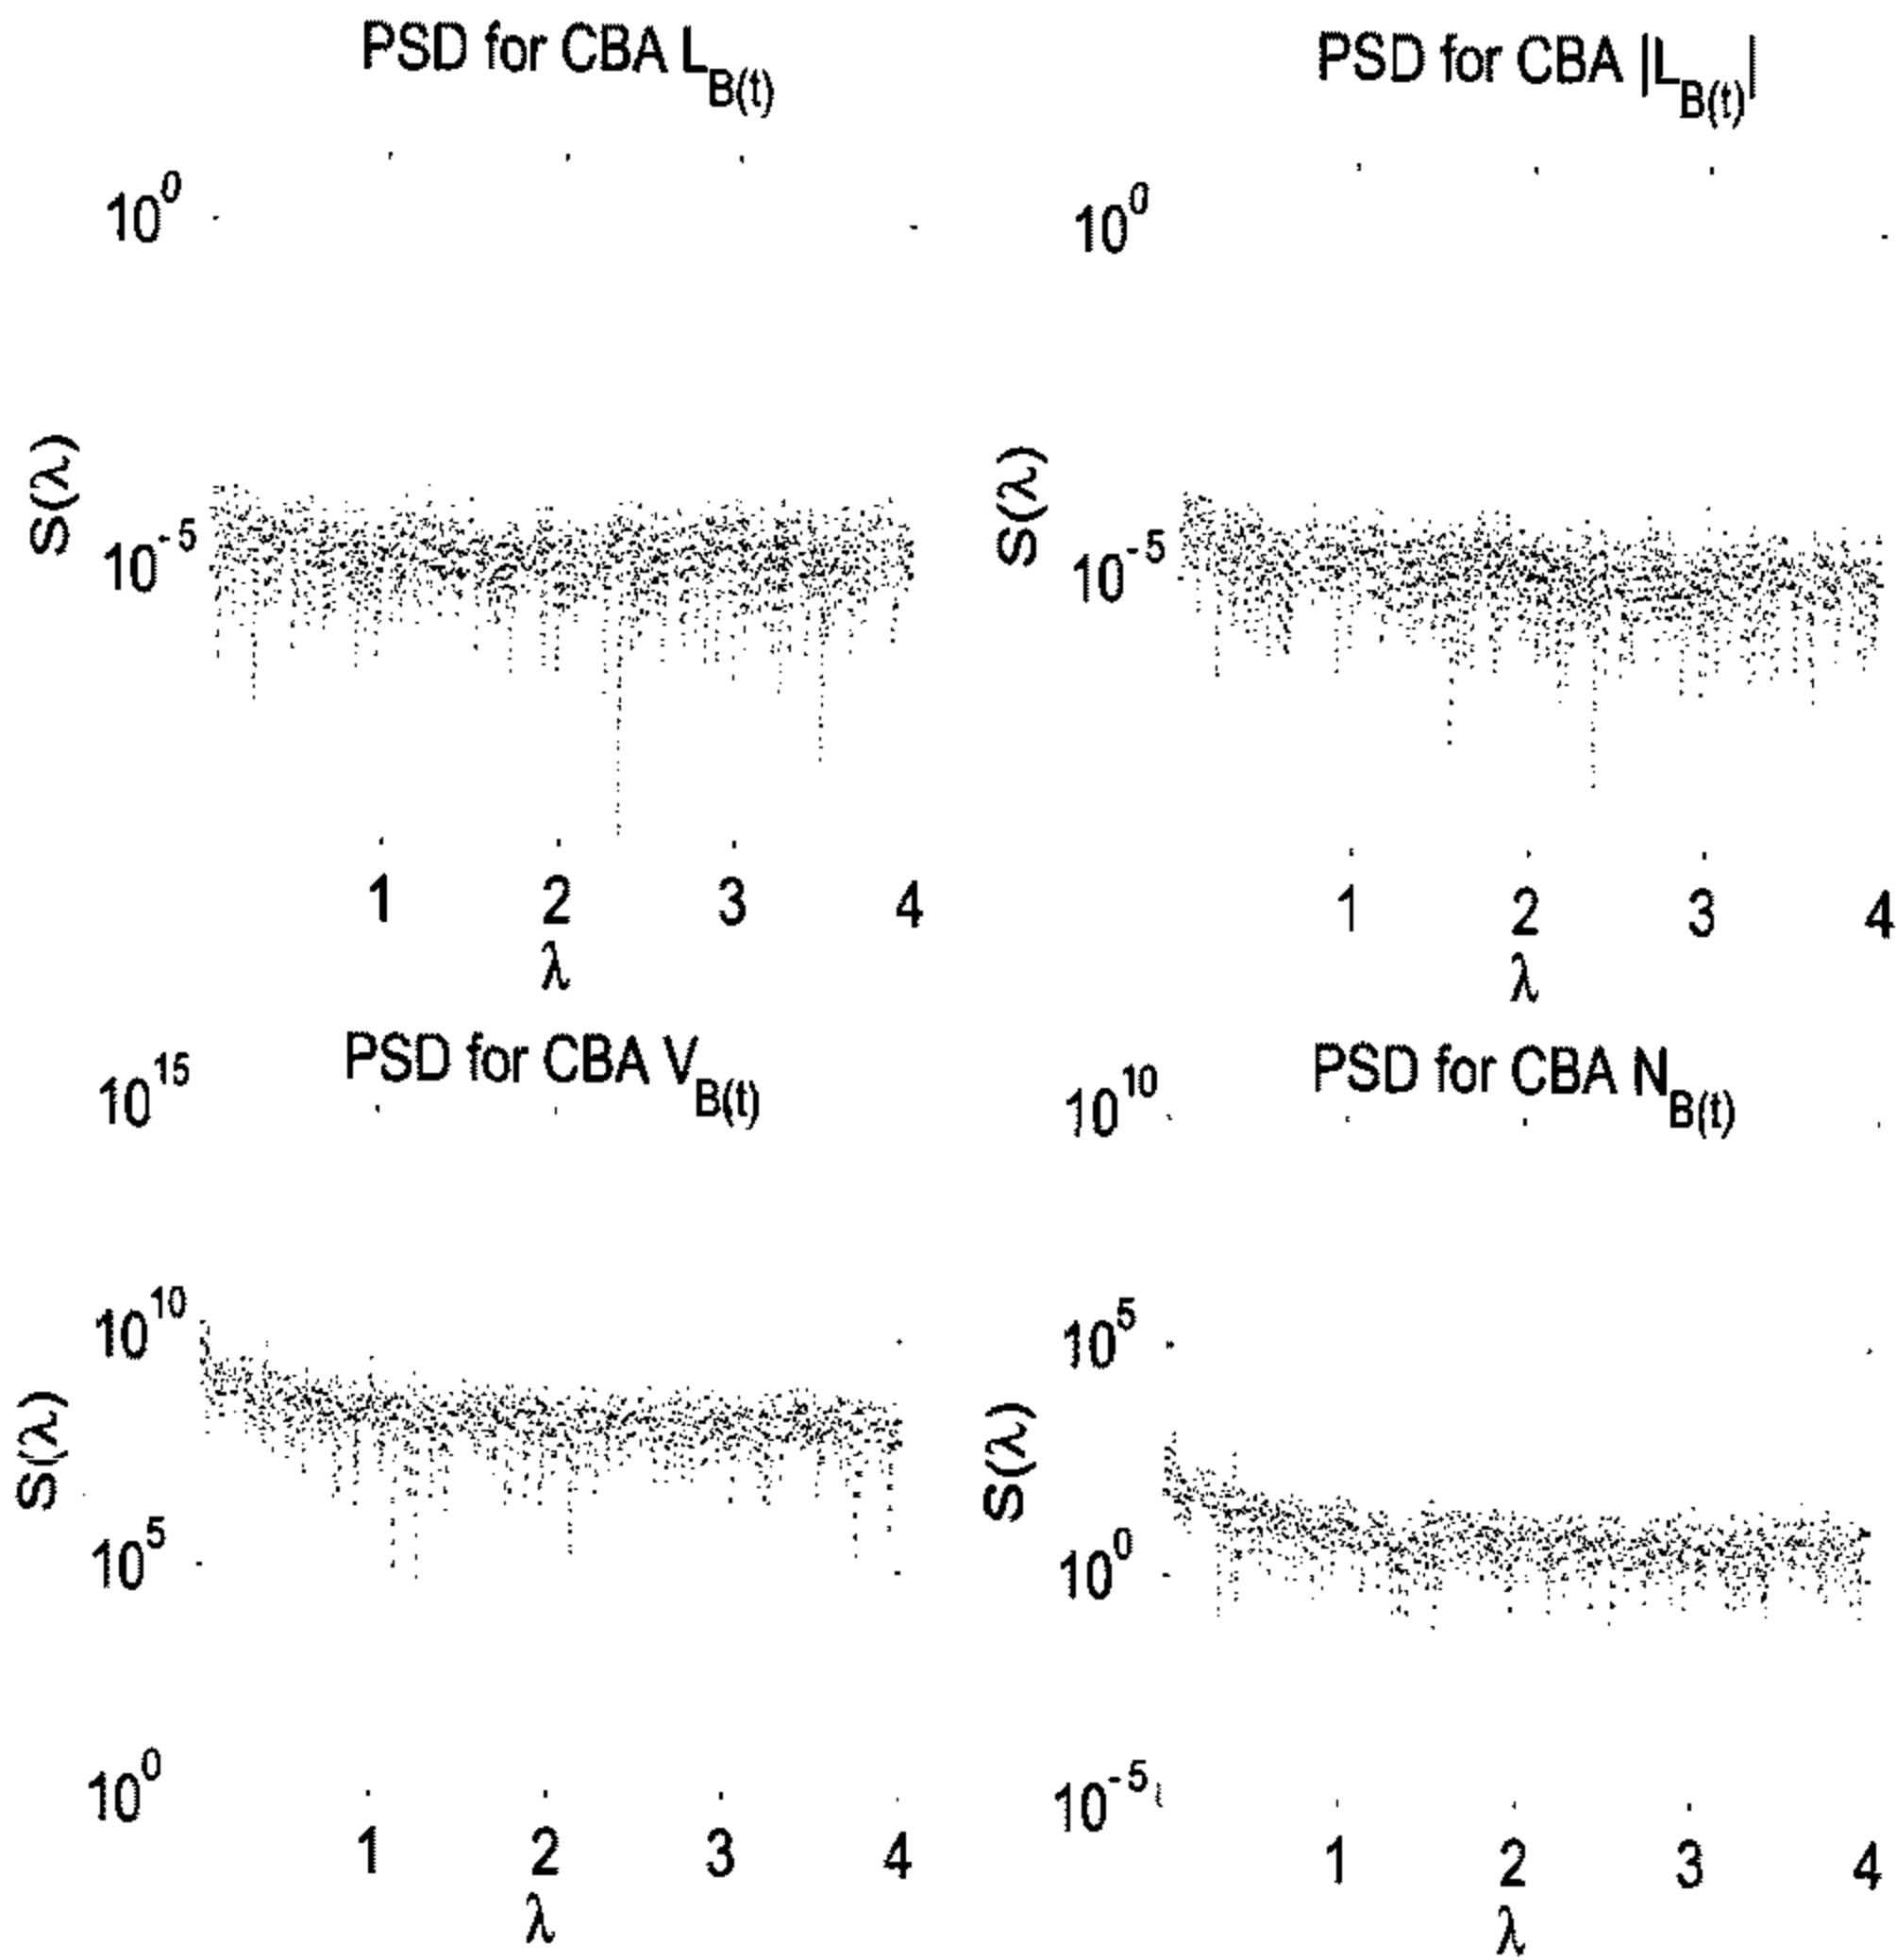


Figure 3.11: Spectral Density of $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ in business time. We can see that the periodic nature of the data has been accounted for.

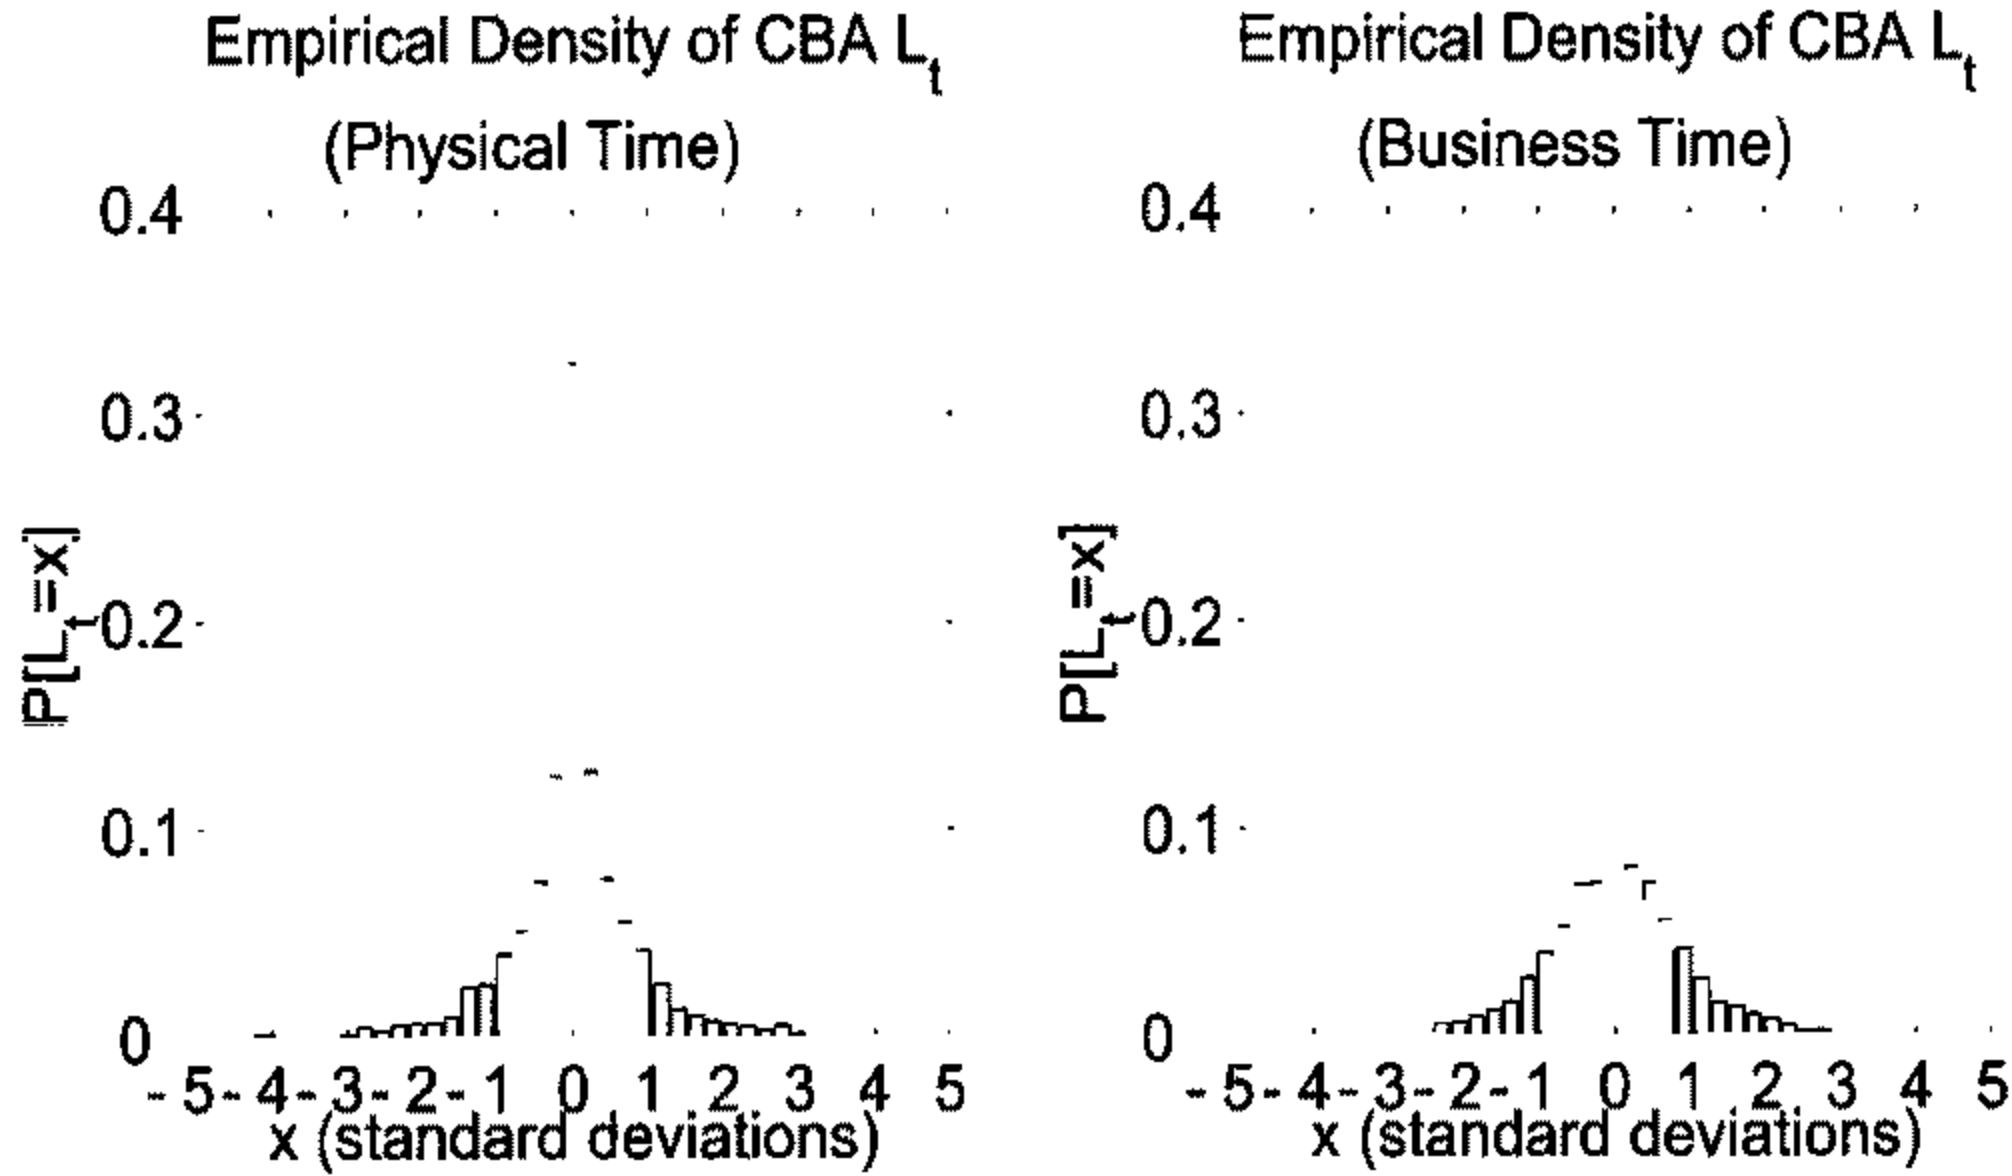


Figure 3.12: The left plot shows the empirical density for CBA returns. Clearly visible is the large central spike at $\mathbf{P}[L_{t,\tau} = 0]$. On the right is shown the empirical density of CBA returns in business time as defined by the master volatility measure.

1 account for less than 1% of the return series, while zero's of type 2 and type 3 make up almost 30% of the total data. Concentrating on the intraday noise process, we can examine the distribution of zero returns over the course of a day, obtaining an estimate of the probability of a zero return at each intraday interval. In Fig. 3.13 the distribution of zero returns is shown and it can be observed that they are strongly related to the intraday activity patterns of Fig. 3.3 and the periodic nature of the intraday noise process. Following the method outlined in Section 3.2.3 we can map the ASX data onto a suitably chosen new time scale. In fact, by choosing the master volatility measure as the basis for our business time, we find that the level of zero return enhancement is made constant across trading hours. Other measures we used to create a business time scale have so far failed to produce the consistent levels of zero returns achieved by using the master volatility measure as shown in Fig. 3.13.

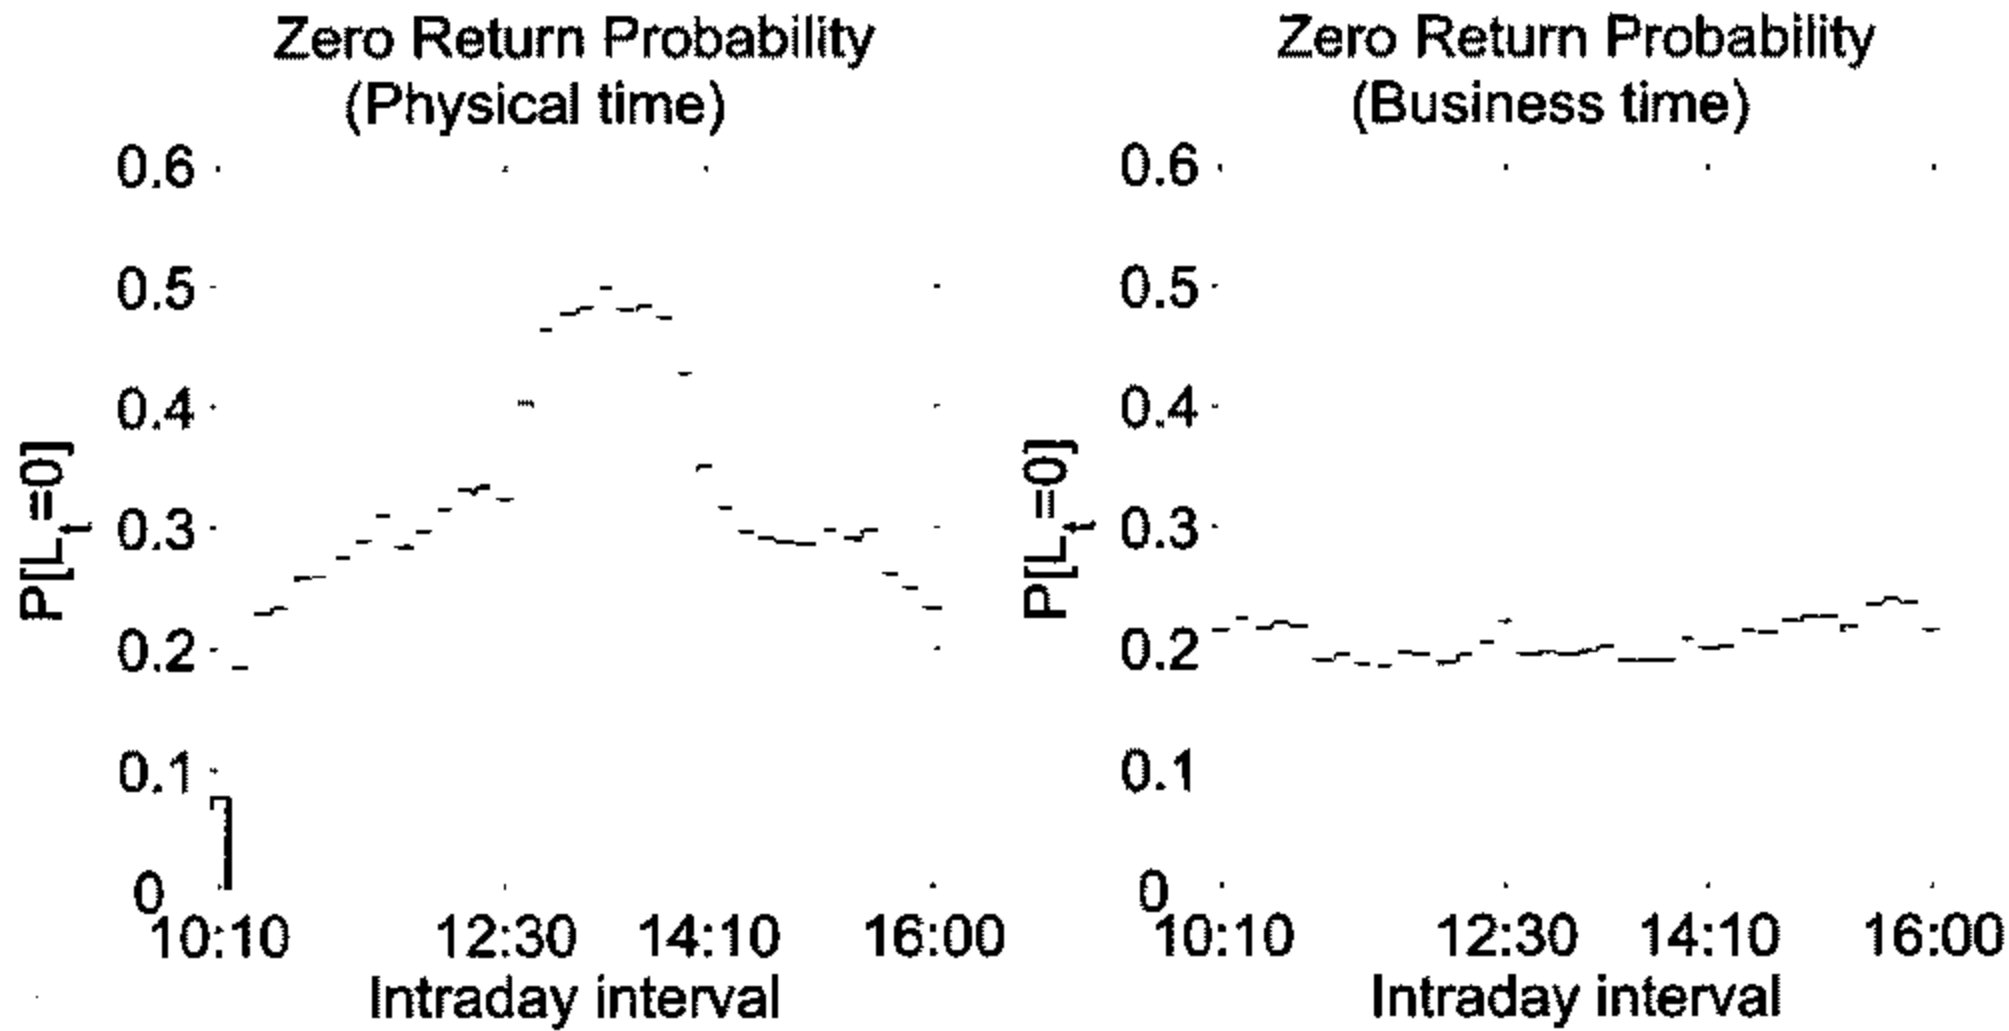


Figure 3.13: The left plot shows the zero return probability, $P_i[L_{t,\tau} = 0]$, calculated at the i^{th} 10 minute intraday interval. The plot on the right shows the zero return probability over 10 minute intervals in business time, constructed using the master volatility measure.

3.3 Distribution properties

The distribution of the log-return is possibly the most studied quantity in finance. The behaviour of the tails of the distribution are of particular interest as it is these that determine the frequency of market booms and crashes. With high frequency data it is now possible to study the distribution of volume, $V_{t,\tau}$ and transaction frequency, $N_{t,\tau}$ which together with the absolute log-return, $|L_{t,\tau}|$, offer insight into the properties of the market volatility. There is no longer any question about financial time series displaying fat tailed distributions. Instead, debate now rages over the extent of the fatness. There exist two main factions: those backing the stable Paretian hypothesis and those advocating finite variance. In the present study we set out to examine the distribution behaviour of the different time series. We start with a parametric approach where we investigate how well the previously described models of Chapter 2 fit the distributions of high frequency data. We then proceed with a non-parametric approach, examining the empirical density and distribution functions. This approach allows us to examine the data free of the burden of pre-supposition and determine how fat the different distribution tails are.

3.3.1 Parametric estimation

We now use high frequency log-return data to calibrate some of the previously proposed models. We look at the subordinated models examined in Chapter 2 but also include another more recent model, the truncated Lévy flight (TLF) of Mantegna and Stanley (1994). By turning to high frequency data we will be able to gain a better picture of the tail behaviour of the log-return distribution and thereby draw conclusions as to which model offers the best description of the unconditional distribution. We also examine the performance of the conditional GARCH(1,1) model and its ability to describe not only the distribution but also the dependence structure of the data. To fit the models we use the method of maximum likelihood estimation on 10 minute log-return data.

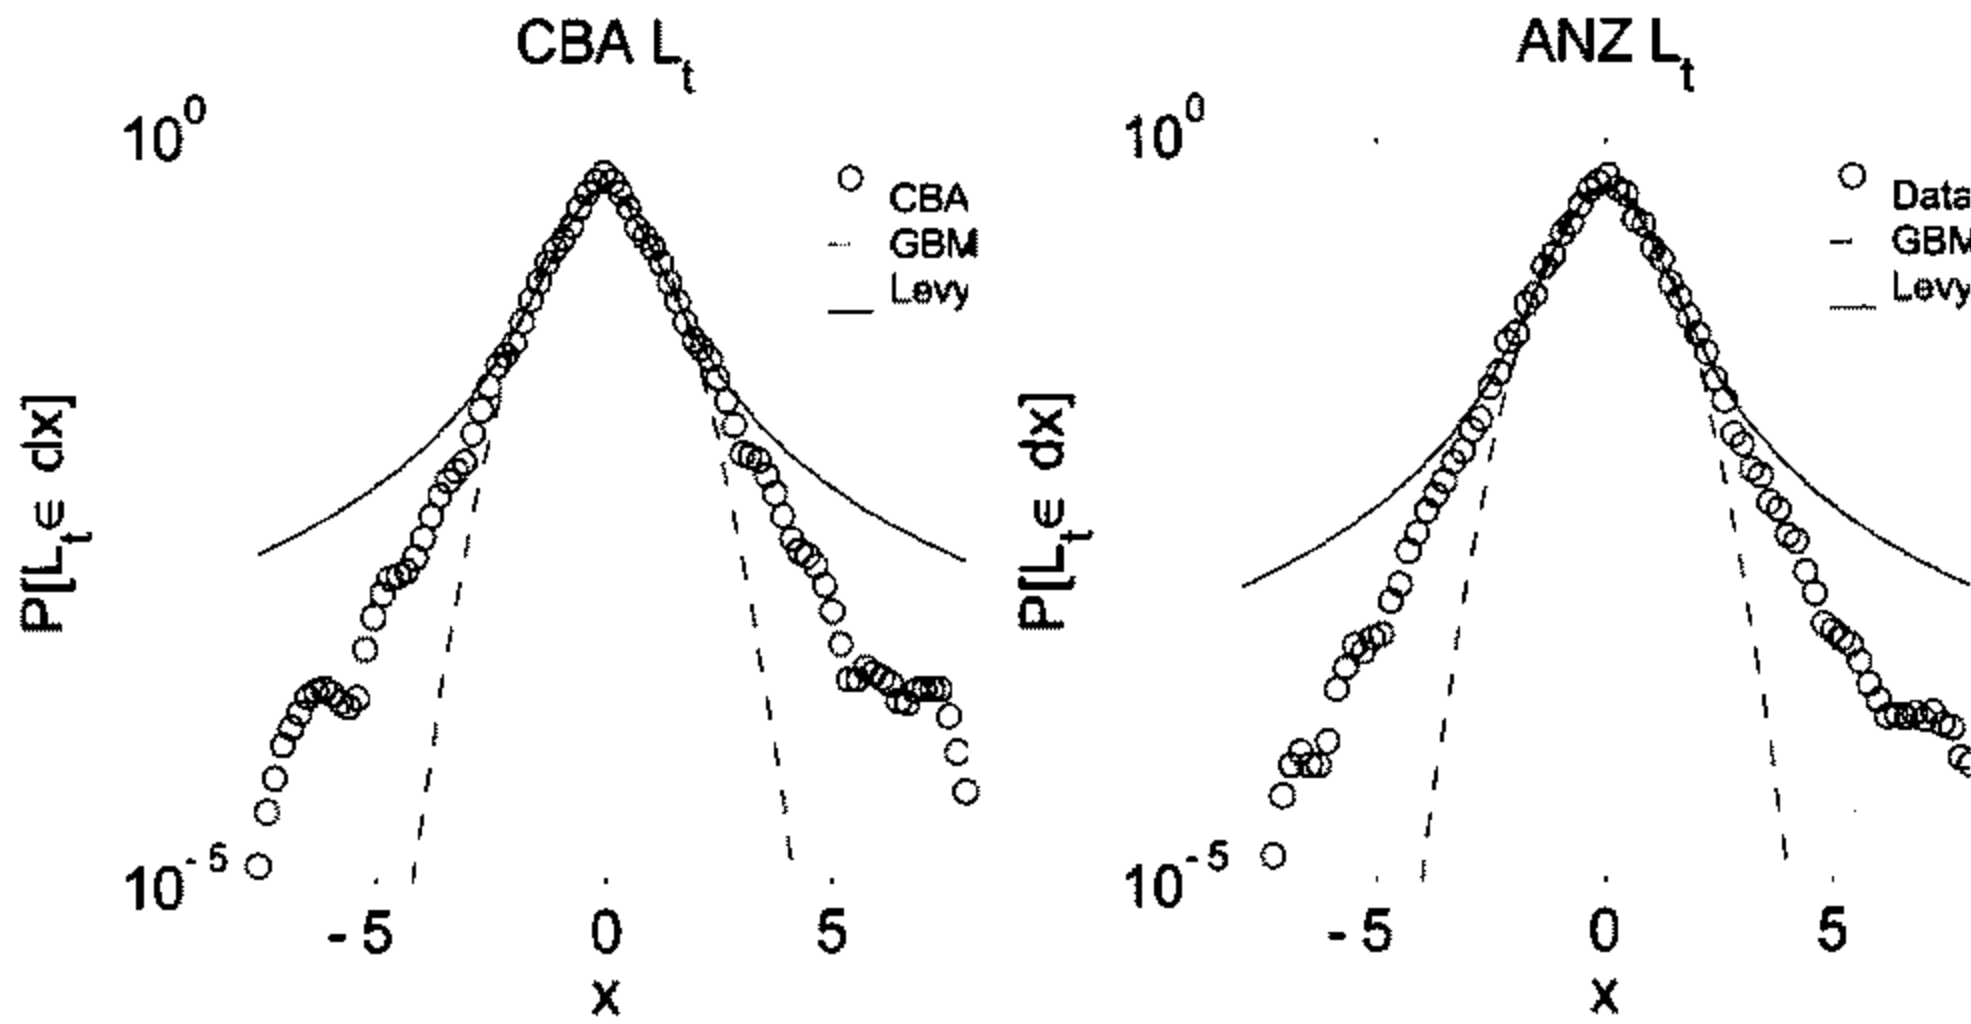


Figure 3.14: Parametric estimation for Lévy distribution using CBA and ANZ high frequency log-return.

The α -stable model

Fig. 3.14 displays the results of fitting the α -stable model to the unconditional log-return distribution via maximum likelihood estimation. The Gaussian estimate is included in each figure in order to provide a comparison with the GBM model.

For both ANZ and CBA stocks we find that the tails of the Lévy distribution overestimate the empirical data. There is a clear departure from data around the scale of 2 to 3 standard deviations. In this case it appears that the infinite variance of the stable Paretian hypothesis is unnecessary, although the distribution does offer somewhat of an improvement over GBM.

The truncated Lévy flight

We have just examined the two extreme cases for modelling the log-return distribution. On one hand the normal distribution function is seen to underestimate the frequency of large fluctuations in the log-return. On the other hand the Lévy distribution, with its infinite variance property, tends to over estimate the data. In order to curtail the extreme fatness of the Lévy distribution, Mantegna and Stanley (1994) proposed truncating the tails of the Lévy distribution at some chosen time scale $\frac{1}{\zeta}$. This idea lead to the proposal of the truncated Lévy flight. The truncated Lévy distribution can be defined by the characteristic function,

$$\phi(\theta) = \exp \left\{ -\sigma \frac{(\zeta^2 + \theta^2)^{\alpha/2} \cos(\alpha \arctan \frac{|\theta|}{\zeta}) - \zeta^\alpha}{\cos(\frac{\pi\alpha}{2})} \right\} \tag{3.9}$$

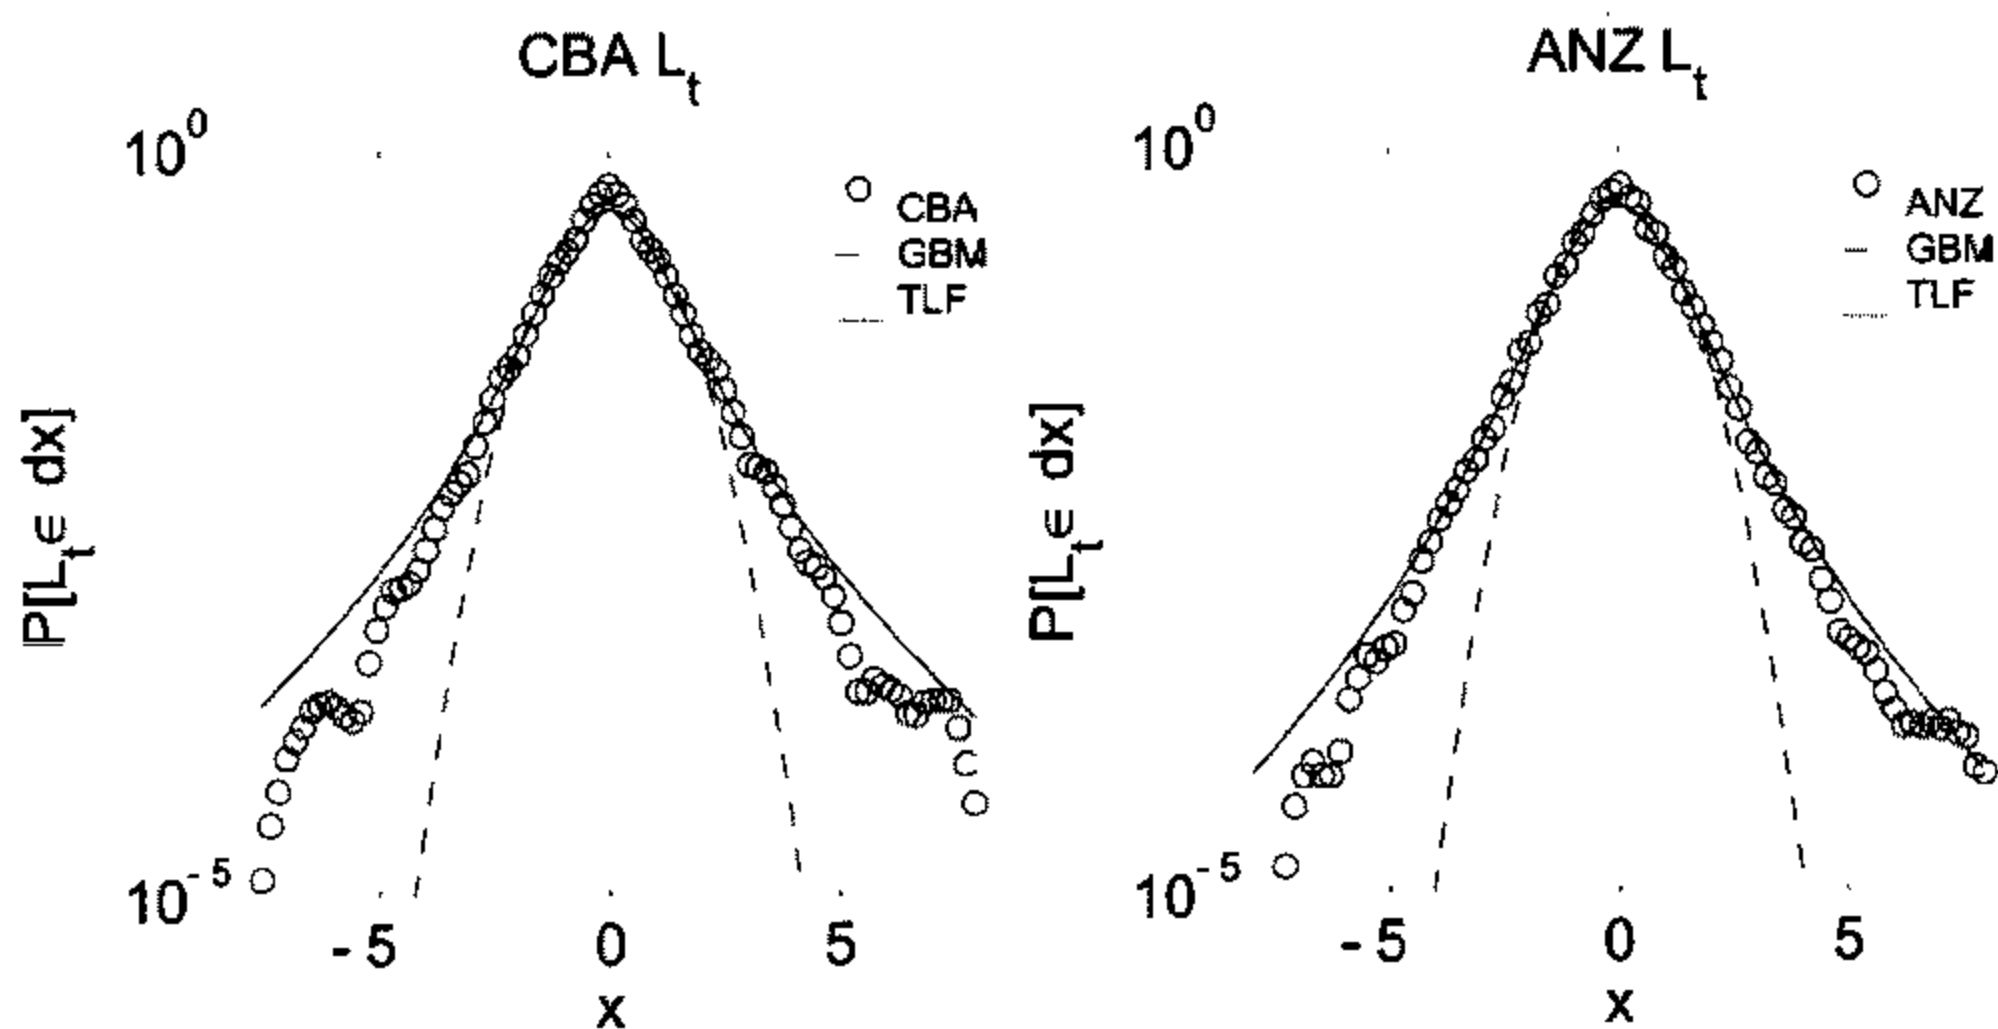


Figure 3.15: Parametric estimation for TLF distribution using CBA and ANZ high frequency log-return.

such that the distribution reduces to a Lévy distribution for $\zeta \rightarrow 0$ and to a Gaussian for $\alpha = 2$,

$$\phi(\theta) \sim \begin{cases} \exp(-\sigma|z|^\alpha) & \text{for } \zeta \rightarrow 0 \\ \exp(-\sigma|z|^2) & \text{for } \alpha = 2. \end{cases} \tag{3.10}$$

For positive scale parameter $\zeta > 0$ the distribution has finite variance, and therefore the central limit theorem guarantees that the truncated Lévy distribution converges towards Gaussian. Mantegna and Stanley (1994) and Bouchaud and Potters (1997) used the TLF to investigate high frequency S&P500 index data.

Fig. 3.15 displays the truncated Lévy distribution estimates for CBA and ANZ data. Unlike the results of Mantegna and Stanley (1994) and Bouchaud and Potters (1997), we find that the truncated Lévy distribution still overestimates Australian equity data in the extreme tails of the distribution. Also we find that the model tends to describe the right tail of the distribution better as the left tail seems to decrease more rapidly. This finding indicates one of the many difference between the behaviour of equity log-returns and the behaviour of index log-returns.

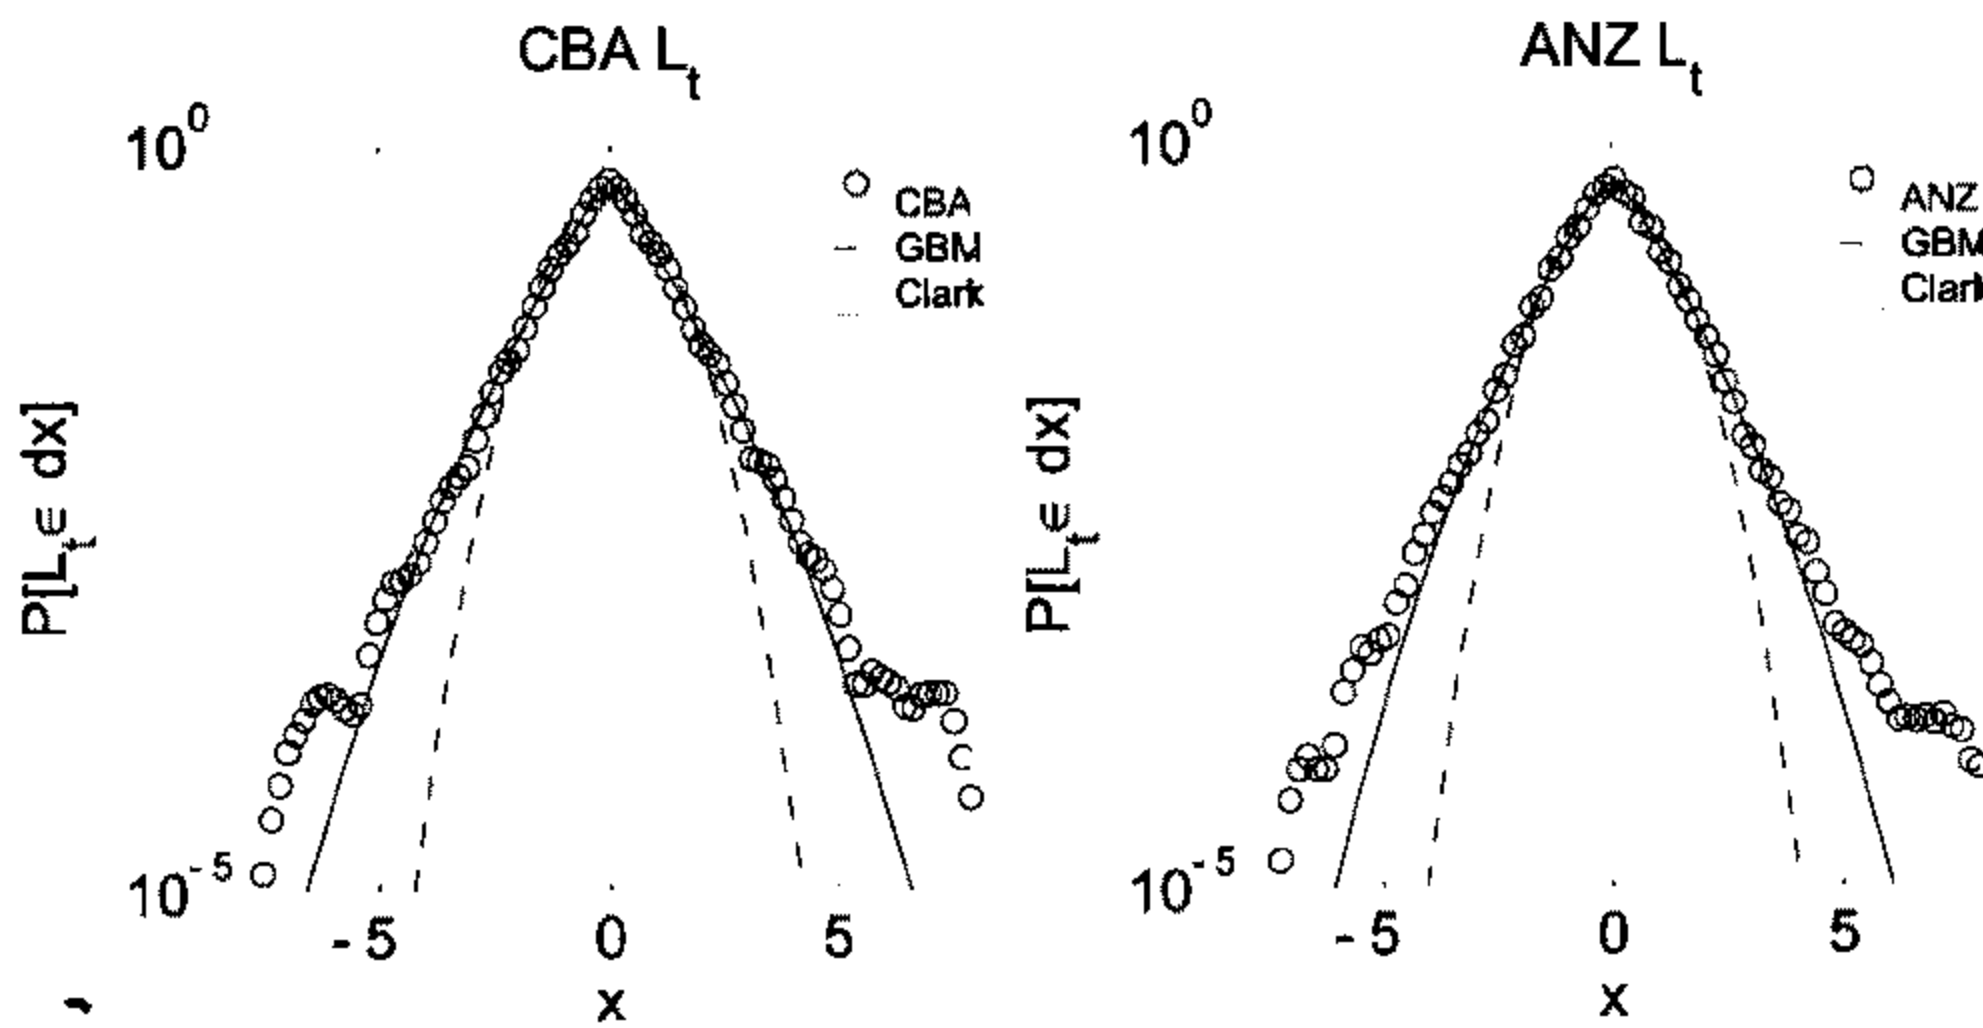


Figure 3.16: Parametric estimation for Clark distribution using CBA and ANZ high frequency log-return.




Figure 3.17: Parametric estimation for Student- t (Praetz) distribution using CBA and ANZ high frequency log-return.

The Clark model

Fitting Clark's lognormal model to the log-return data (Fig. 3.16) we observe that it describes the distribution up to the range of approximately 4 standard deviations. The finite nature of the moments of Clark's model mean that the resulting distribution tends to underestimate the extreme tails of the data. The distribution properties can be seen to be closer to the data than those of the Gaussian. In the case of CBA, the Clark model follows the data until 5 to 6 standard deviations.

The Student- t model

In Chapter 2 we found that the Student- t model of Praetz (1972) offered the best description of the daily log-return distribution. In Fig. 3.17 we present the estimates for the model on the 10 minute log-return. Unlike the previous findings of Section 2.3.6 where daily data was used, The Student- t model does not offer as good a description of high frequency data. Like the Lévy distribution, the Student- t distribution tends to over estimate the tail thickness. Deviation is seen to occur after approximately 3 to 4 standard deviations. For high frequency data the degrees of

freedom parameter is found to be much lower, $\nu \approx 2$, than for daily data.




Figure 3.18: Parametric estimation for Variance Gamma model using CBA and ANZ high frequency log-return.

he variance gamma model

Fig. 3.18 shows the fitted distribution for the VG model of Madan and Seneta (1990). We see in these figures that the VG model offers the best description of the data of all the models. As for the other models, the VG distribution also cannot match the apparent asymmetry of the data. The right side of the distribution function deviates from the extreme tail behaviour of the data.

Parametric results

Comparing the results of this section to those in Chapter 2, it is easy to see the insight gained by working with high frequency data. Low frequency daily data do not allow for a clear picture of the distribution behaviour, whereas now we can draw conclusions as to the best model available. Tables B.3 to B.8 give the parameter estimates for the unconditional models in this section over a variety of stocks in the data set. In Table B.9 we list the log-likelihood values obtained for each

of the estimates. It is clear that in the case of high frequency data, the log-likelihood values indicate that the variance gamma model of Madan and Seneta (1990) provides the best description of the log-return distribution. Table B.9 also indicates that the truncated Lévy distribution and the Clark lognormal model follow closely behind in terms of performance. The Student- t model ranks fourth, outperforming only the α -stable model and geometric Brownian motion which come in last. We must also note that while the models improve tail description, none of them is able to describe the sharp peak in the central region of the empirical density function.

GARCH

In Fig. 3.19 we show the results of fitting the GARCH(1,1) model to ANZ and CBA high frequency data. Here we see that the discrete time model is able to generate a somewhat better fit of the distribution function for the log-return than the unconditional Gaussian. Although perhaps it is not as accurate as that of the VG model since the GARCH(1,1) tails appear to underestimate the data. Also the GARCH model is unable to reproduce the central peak of the density function.

The noted property of GARCH is that it includes a dependence structure in its description of the data. The GARCH(1,1) model displays an exponentially decaying autocorrelation function for the squared log-return. In Fig. 3.20 we show the autocorrelation function, $\rho(k)$, of the square log-return for the ASX data and the GARCH(1,1) estimate. In both of these figures it is found that the GARCH autocorrelation function decays much faster than that of the equity data. This slow decay in the autocorrelation of the squared log-return indicates a dependence that cannot be described by the GARCH(1,1) model.

3.3.2 Empirical distribution function

There exist several methods to estimate and smooth the density function of a time series, as examined by Silverman (1986). Throughout this chapter we adopt the method of kernel estimation to construct the empirical density function. This approach differs slightly from the work of Bertram (2004), where a simple histogram was used, as the kernel estimate produces a smoother density function.

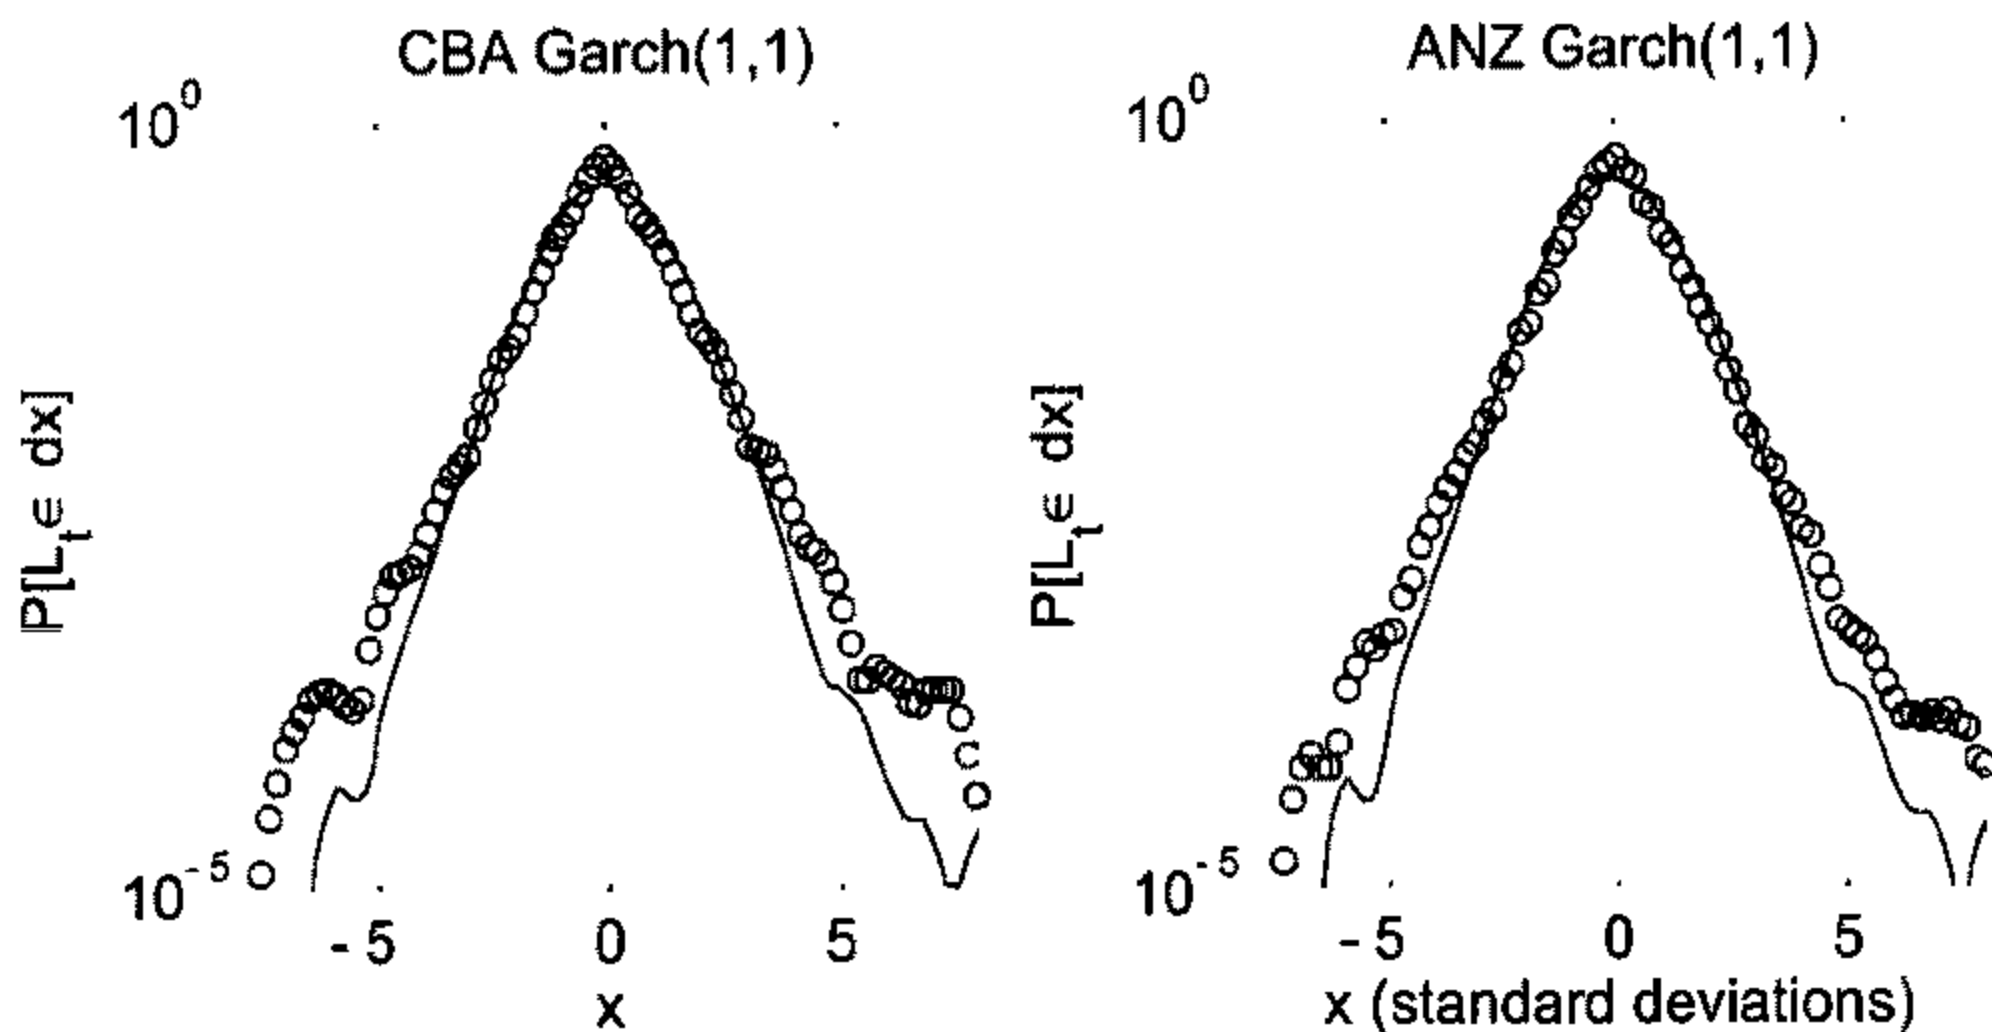


Figure 3.19: Estimation for GARCH(1,1) using CBA and ANZ high frequency log-return.

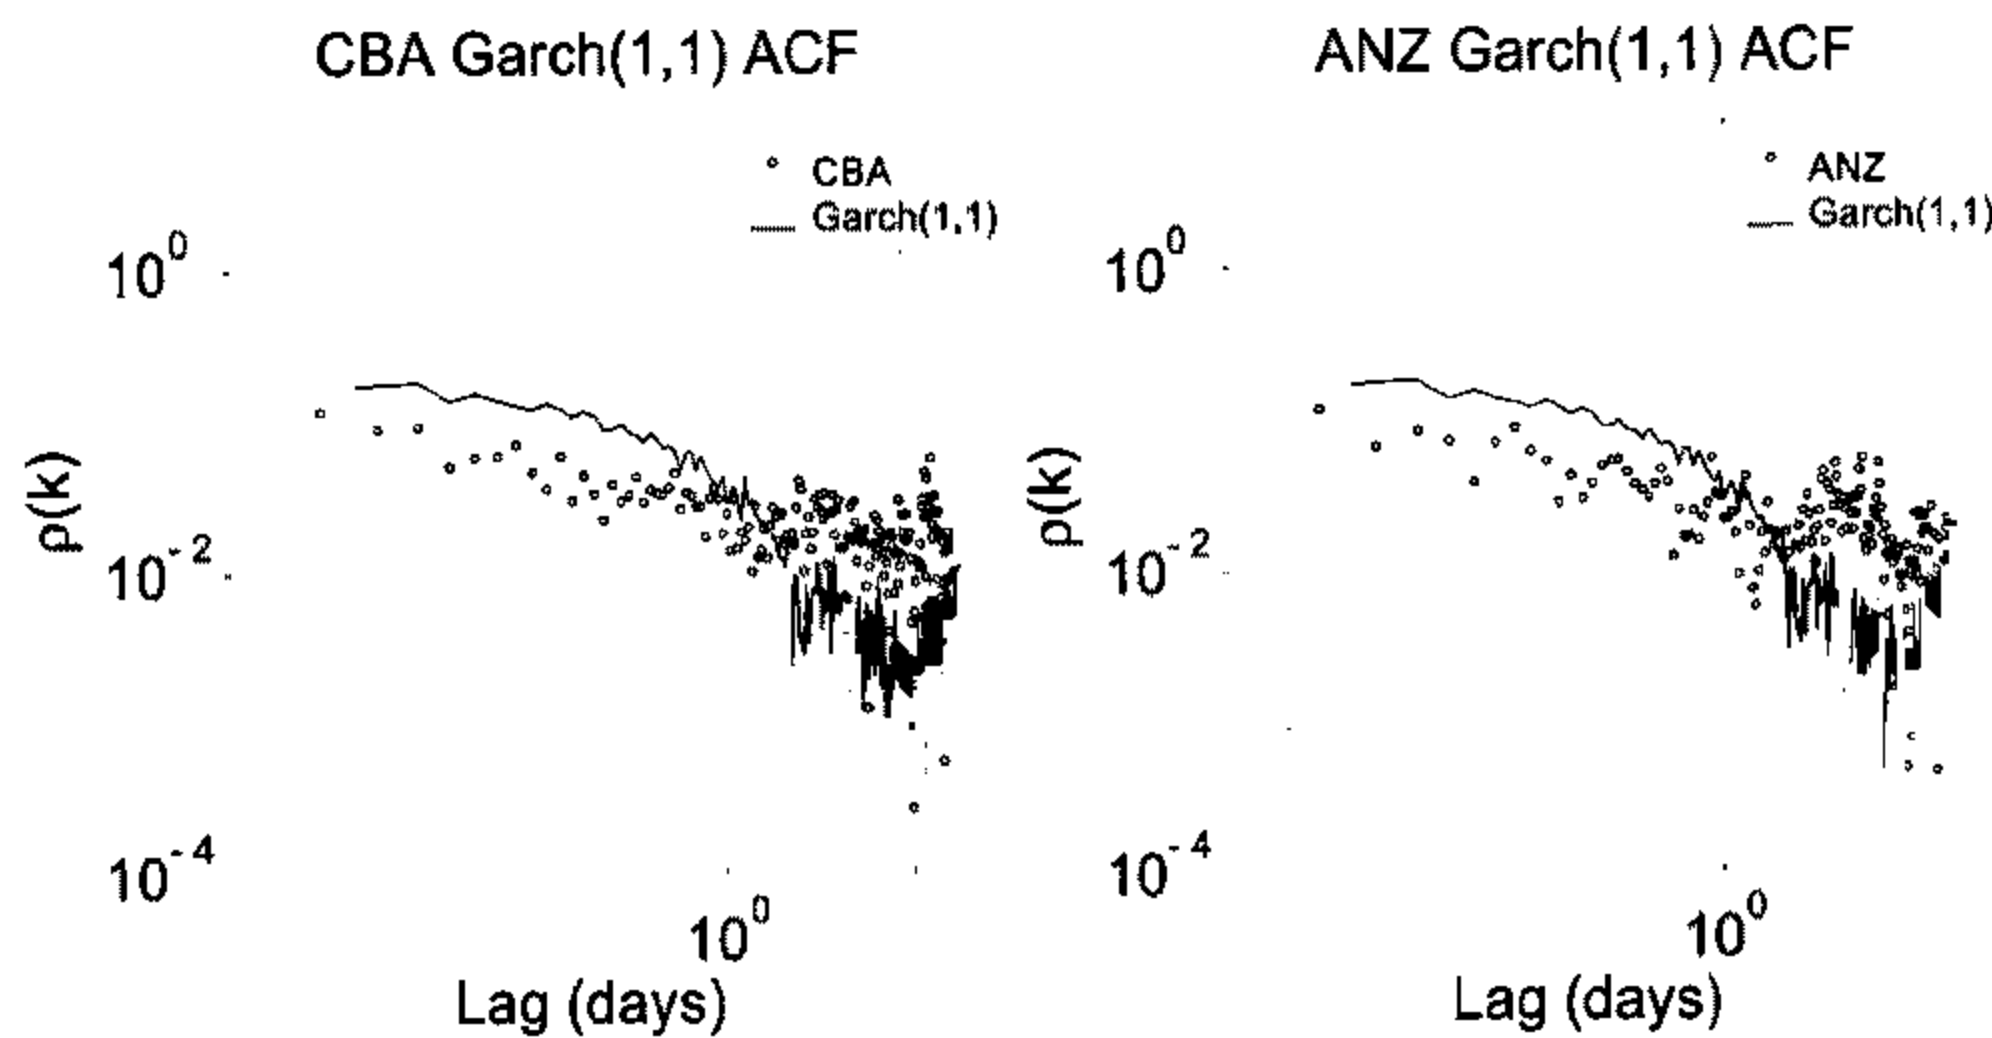


Figure 3.20: Auto correlation function for squared residuals for GARCH(1,1) and high frequency data. We can see here that the GARCH(1,1) ACF decays too rapidly to represent the real data.

Consider a time series, X_1, X_2, \dots, X_n , where n is the number of samples. The kernel estimator for the density function is defined as

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - X_i}{h}\right) \tag{3.11}$$

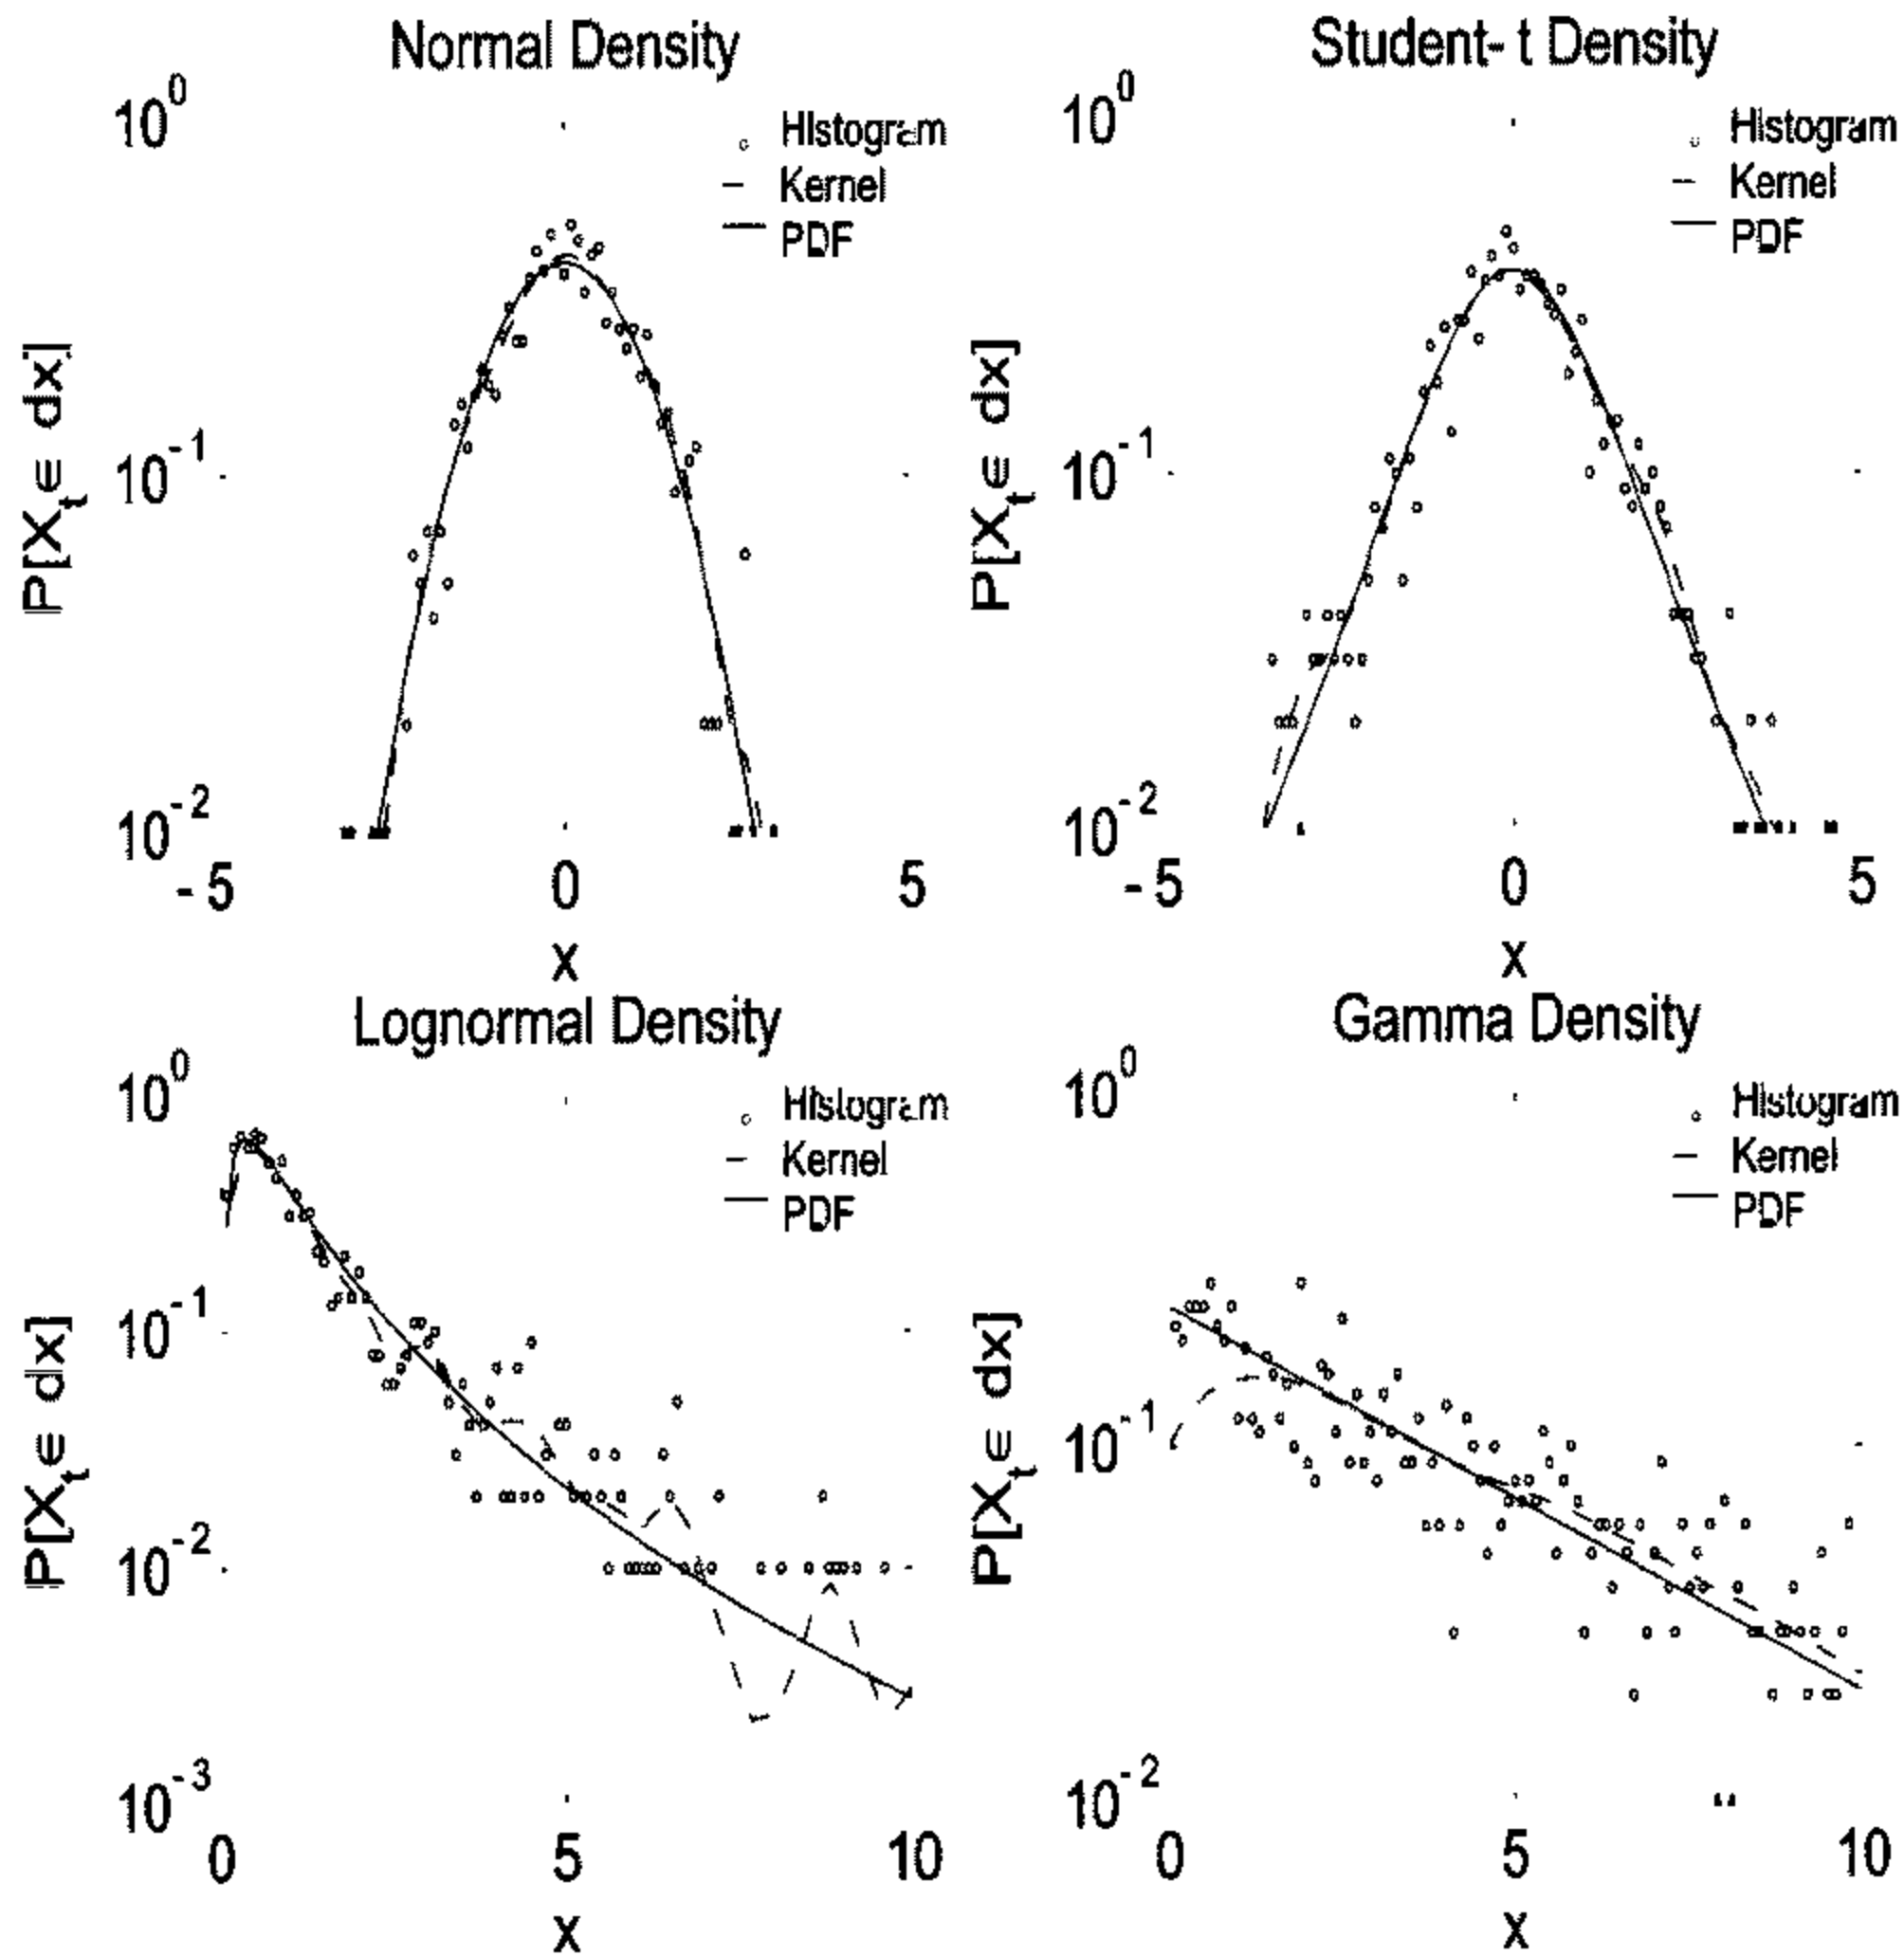


Figure 3.21: Shown are the empirical density functions obtained via histogram and kernel estimates. For each distribution a series of 1000 random variables was used to obtain the estimate. Also shown in this figure are the theoretical density functions.

where K is a kernel function with window width h . Hence the estimate can be thought of as a superposition of ‘bumps’ where K determines the shape of the bumps and h determines their width. Fig. 3.21 displays the density functions of four different time series, each drawn from a different distribution. Shown in each plot is the theoretical density, its histogram estimate and its kernel estimate. For the kernel we use a triangular window with window width chosen as optimal for the

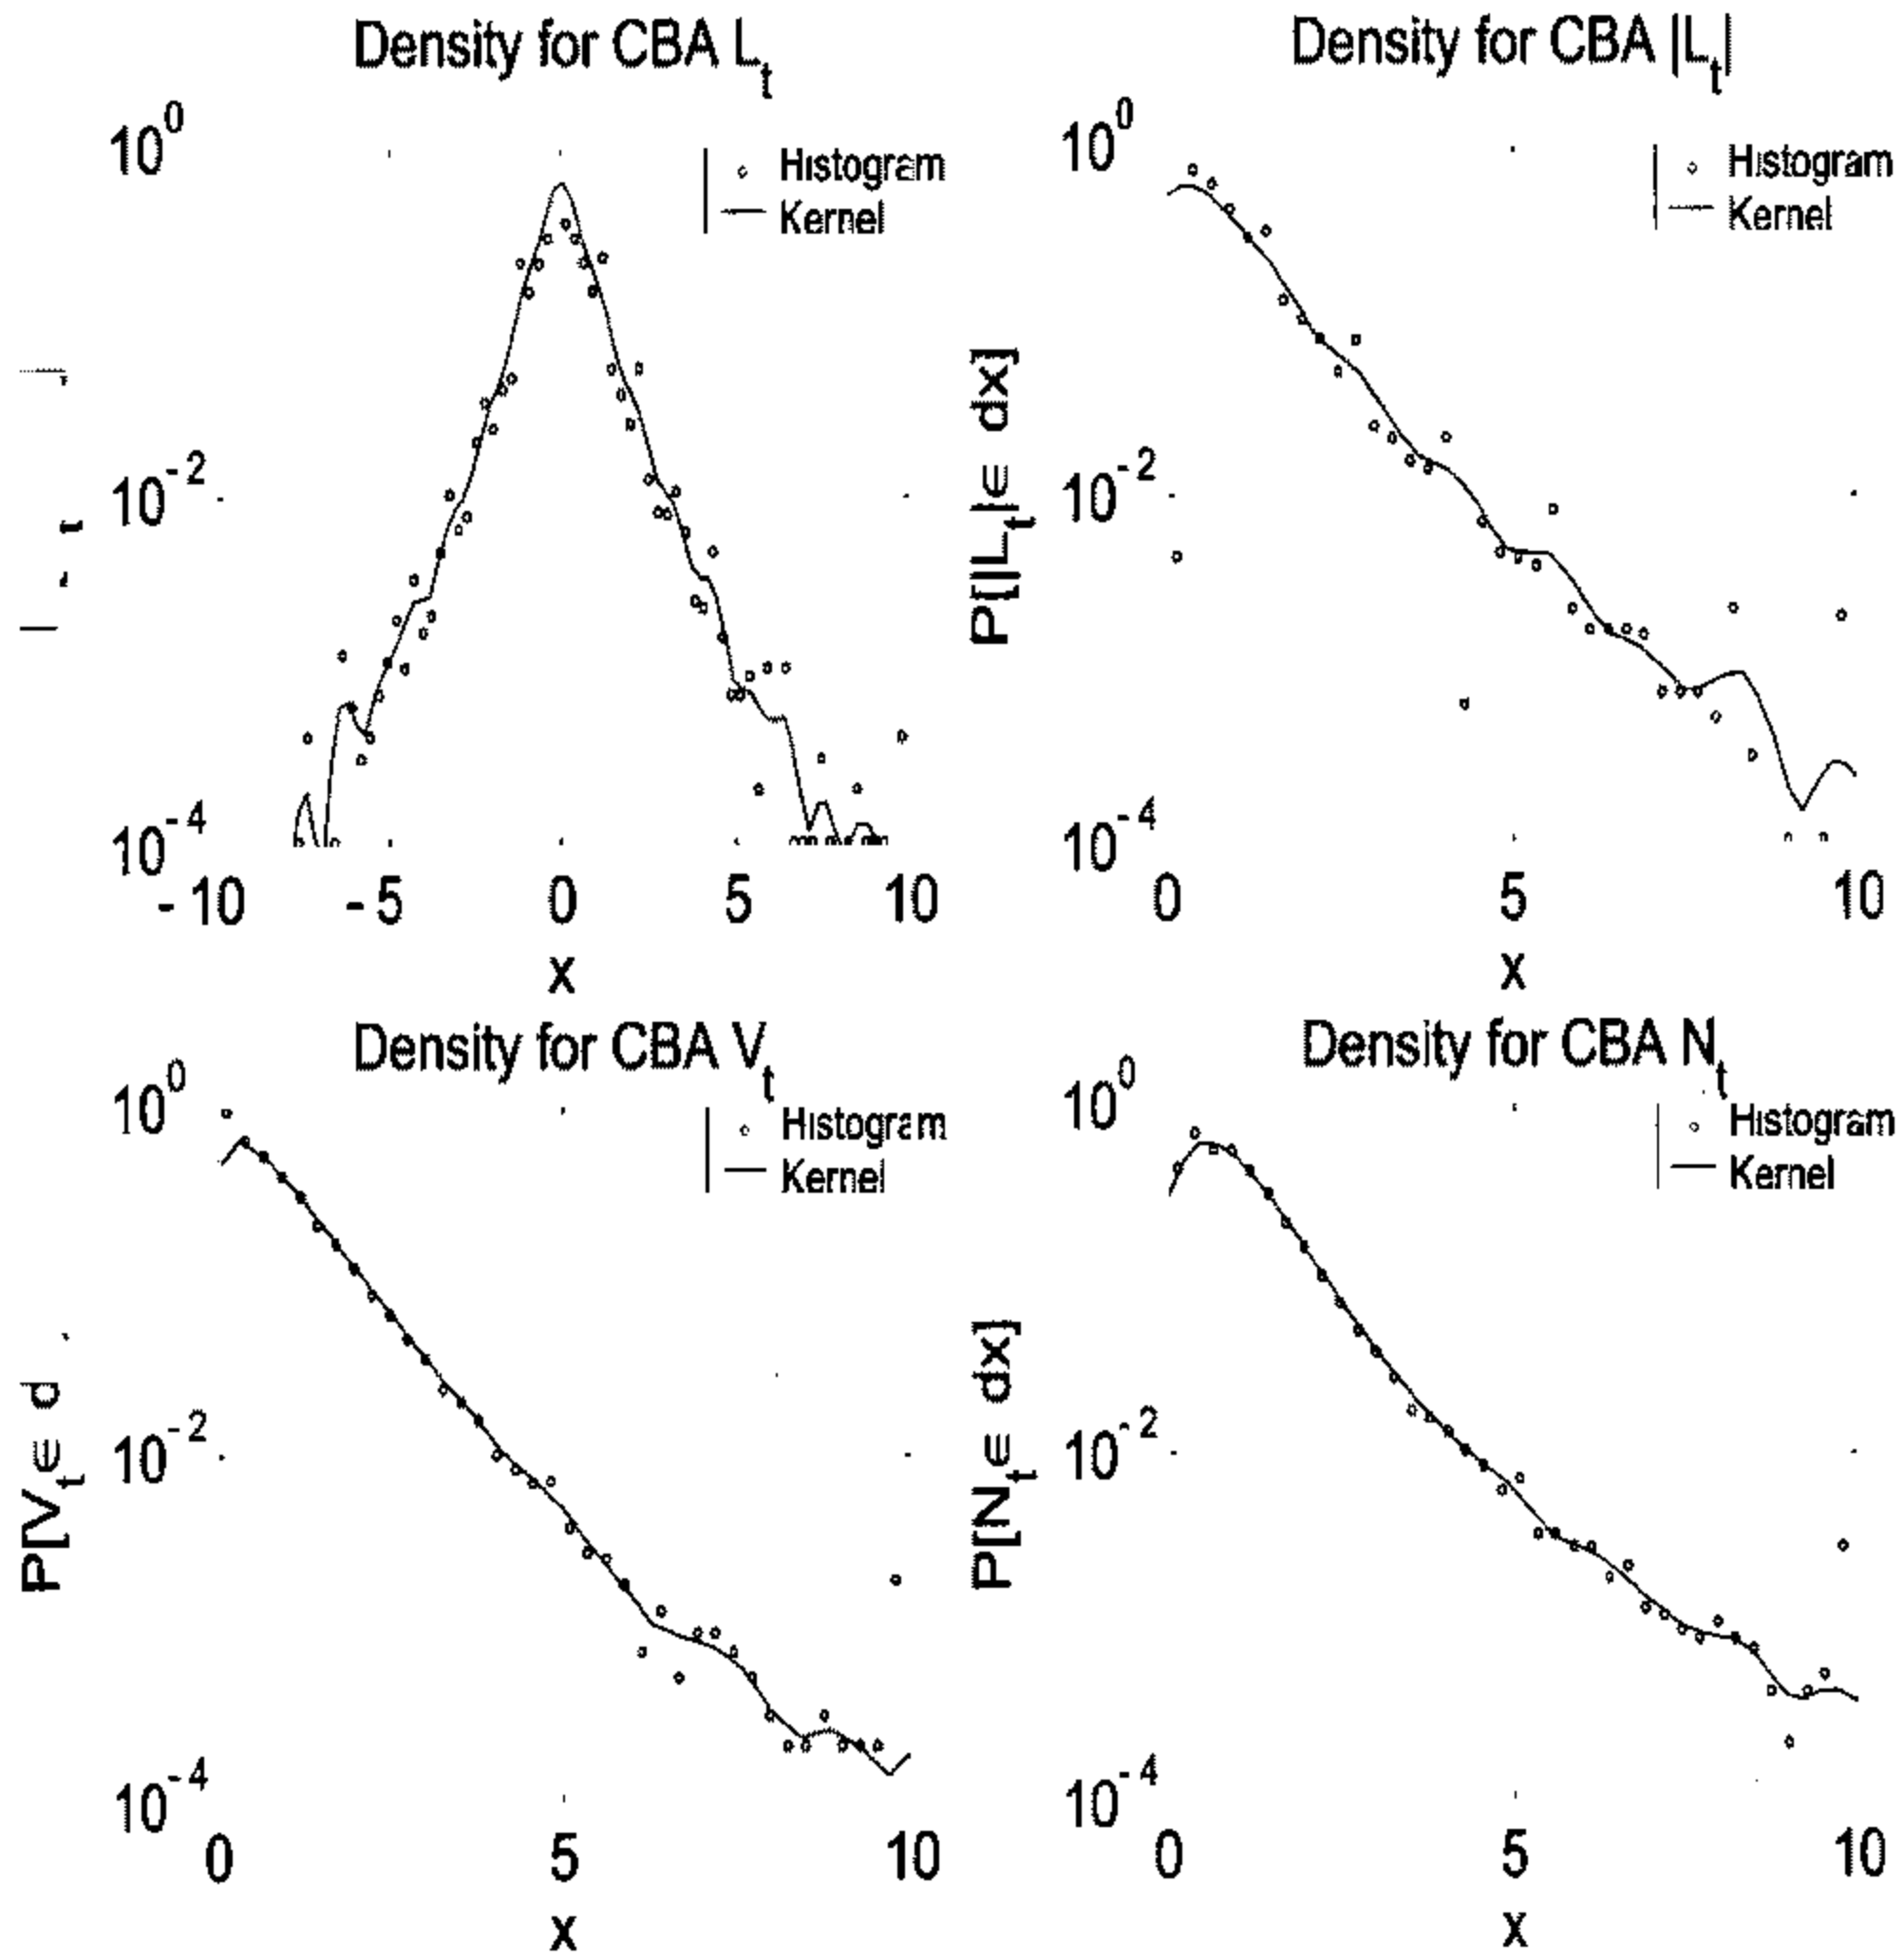


Figure 3.22: Empirical densities for CBA $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. Each series has been normalised to have standard deviation equal to one. Both the histogram and Kernel estimates are shown.

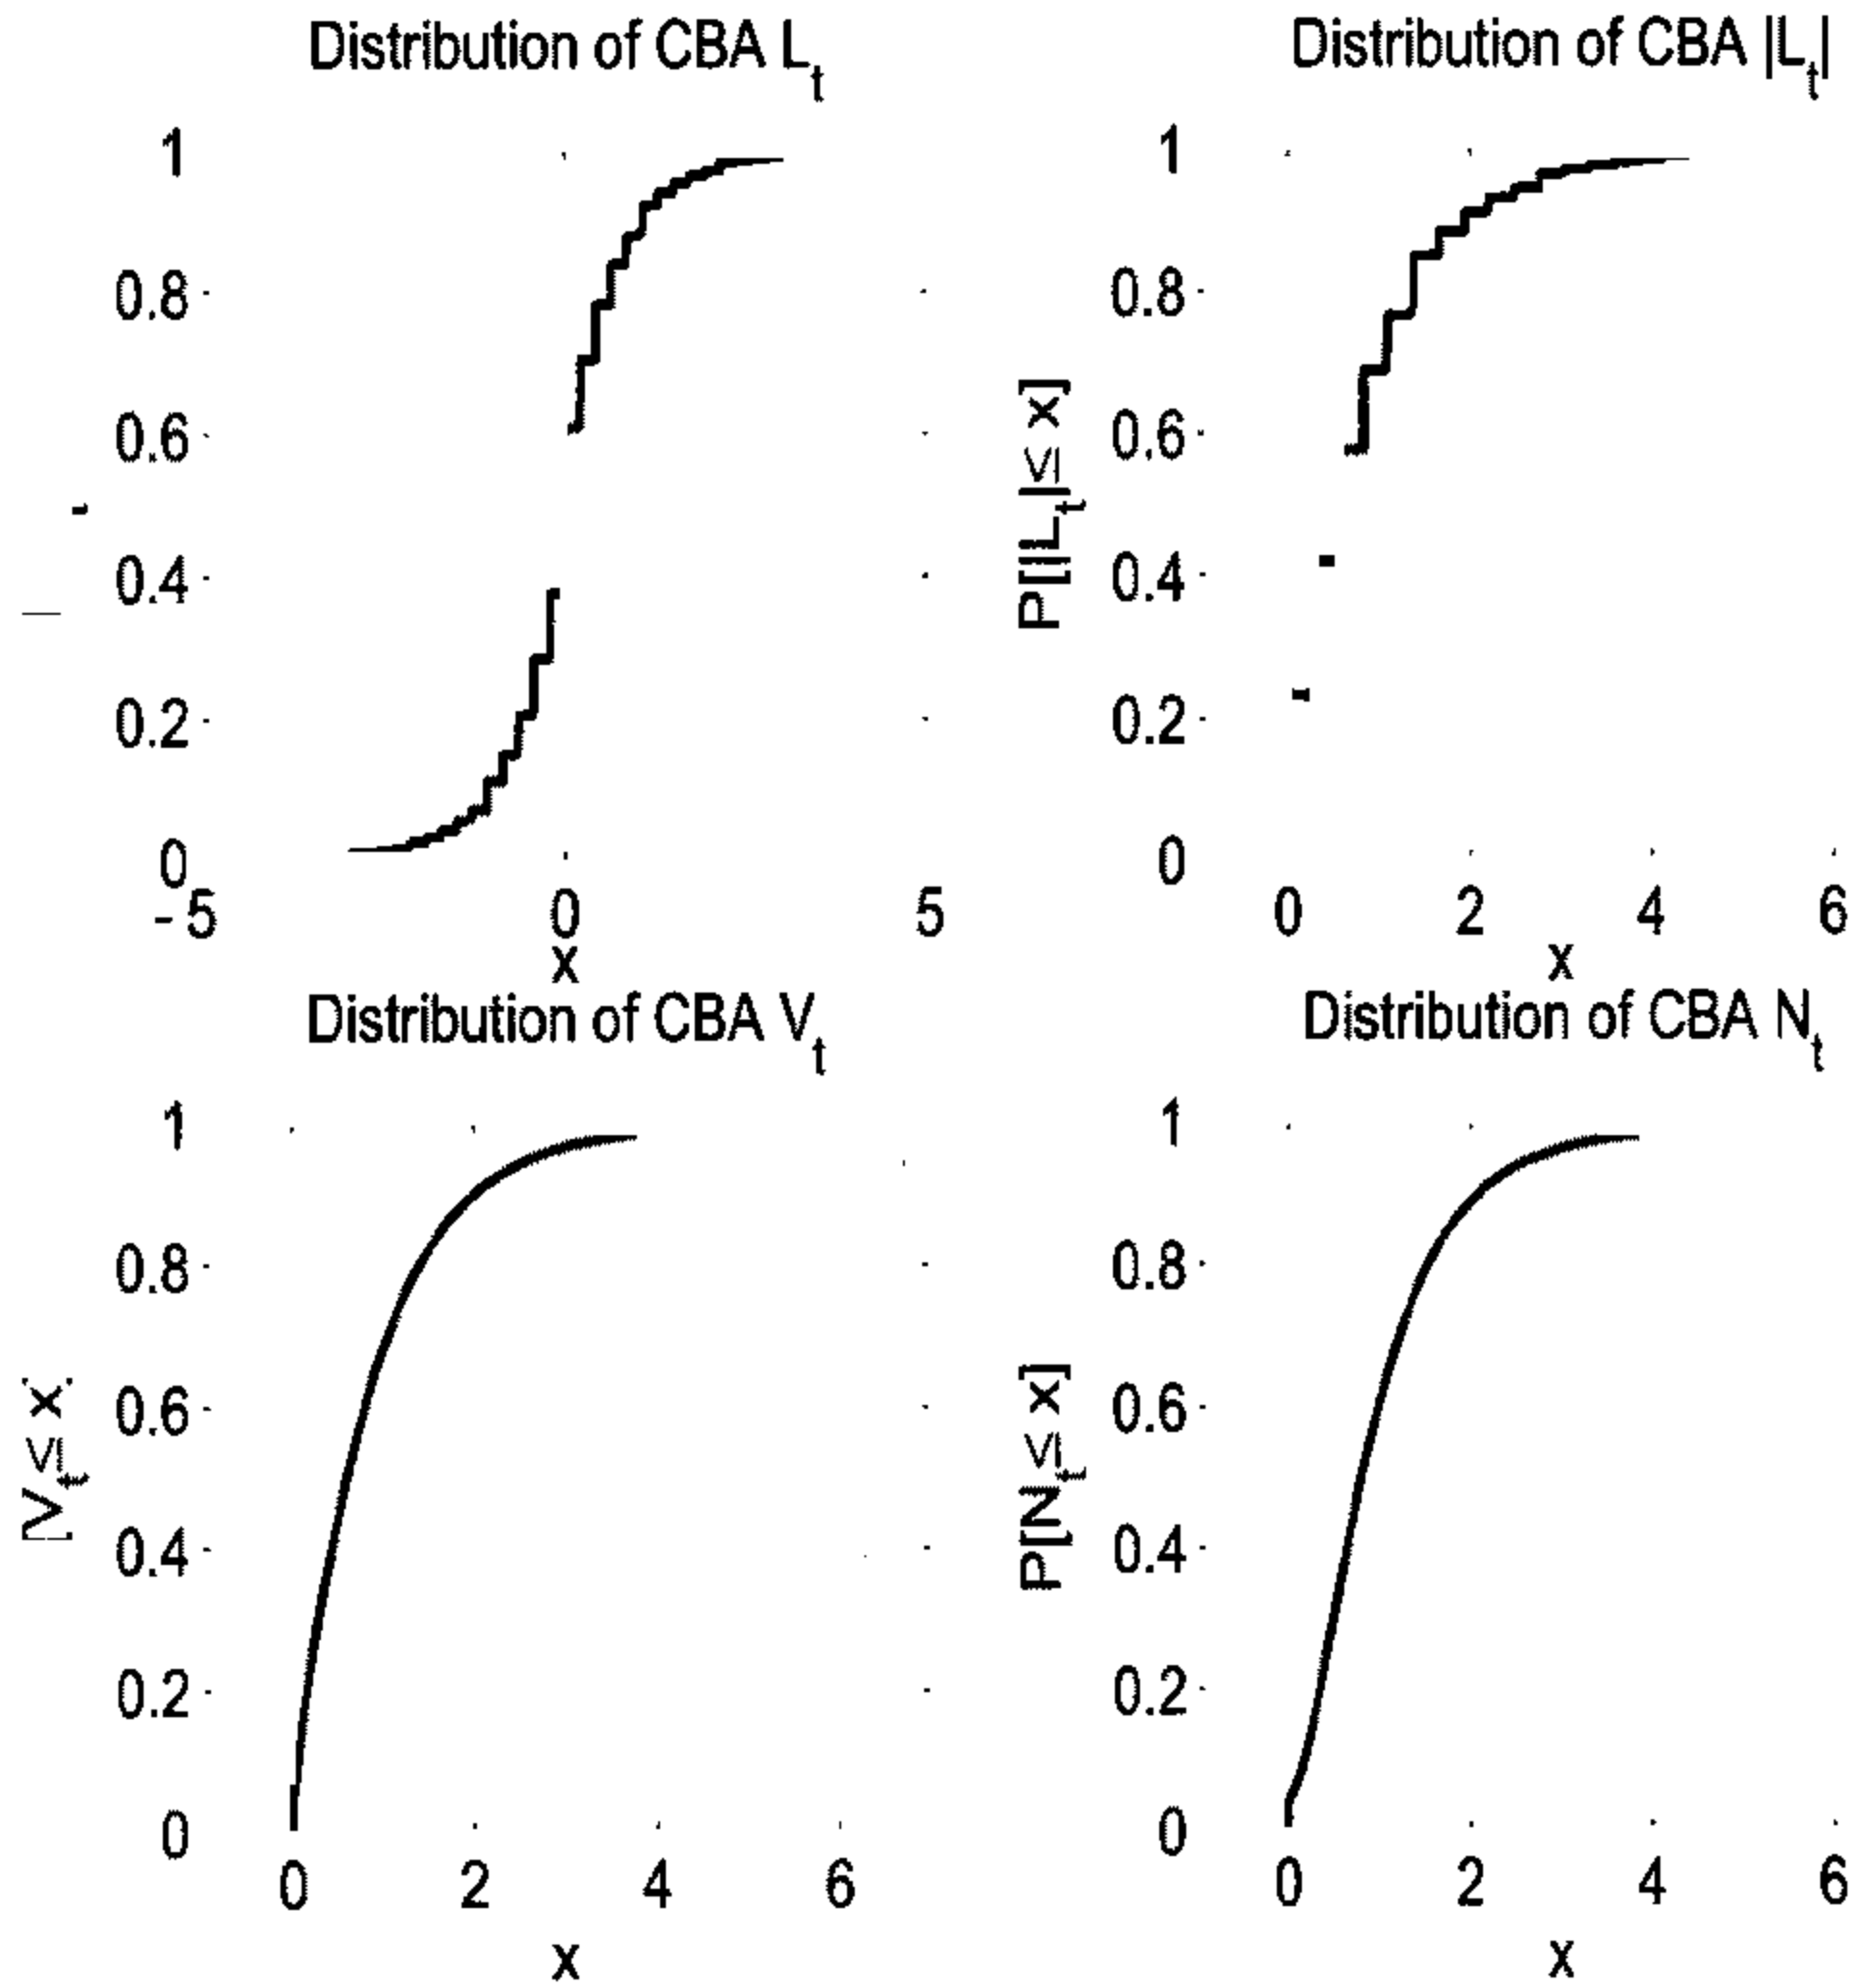


Figure 3.23: The empirical distribution for CBA $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. Each series has been normalised to have standard deviation equal to one. Apparent in this picture is the discrete nature of the log-return and absolute log-return for small relative price changes.

normal distribution. We can see from these plots that the kernel estimate provides a better representation of the density function than the histogram offers.

In Fig. 3.22 we show the density estimates for CBA $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$. Here both the histogram and kernel estimates are shown. While the two estimates offer the same description of the density function, the kernel estimate provides a smoother output that is more useful for observing the properties of the data.

To estimate the cumulative distribution function we can simply integrate the empirical density function. In this study we have made use of the MATLAB statistics toolbox function `ecdf.m` to provide the empirical distribution function together with 99% confidence intervals. This function calculates the Kaplan-Meier estimate for the distribution and uses Greenwood's formula to calculate confidence intervals. The empirical distribution function for CBA $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ is shown in Fig. 3.23. We see in this figure that the return, $L_{t,\tau}$ and absolute return, $|L_{t,\tau}|$ display quite coarse behaviour when compared to the volume, $V_{t,\tau}$ and transaction frequency, $N_{t,\tau}$. This behaviour results from the fact that price changes are limited to cents (*i.e.* two decimal places). The empirical CDF reflects the discrete nature of the log-return for small changes in price. Although the volume and transaction frequency are limited to integer values, these series do not show the same discrete behaviour. The effects of zero return enhancement can also be seen in the CDF of Fig 3.23. While the log-return displays a large probability for $\mathbf{P}[L_{t,\tau} = 0]$, the transaction frequency displays only a small probability for no trade taking place, $\mathbf{P}[N_{t,\tau} = 0]$.

3.3.3 Power law behaviour of equity data

Power laws have a long history in economic modelling. Recall from Chapter 2 that in 1897, Vilfredo Pareto characterised the distribution of household income by a power law, *i.e.*

$$\mathbf{P}[I > x] \sim x^{-\alpha}. \quad (3.12)$$

Such a distribution will have a finite moment of order k , only if $k < \alpha$. As we have seen in Chapter 2, Mandelbrot's stable Paretian hypothesis introduced power law behaviour to asset price modelling with $0 < \alpha < 2$. More recently the works of Amaral et.al (1997), Buldyrev et.al. (1997), Gopikrishnan et.al. (1998) and Gabaix

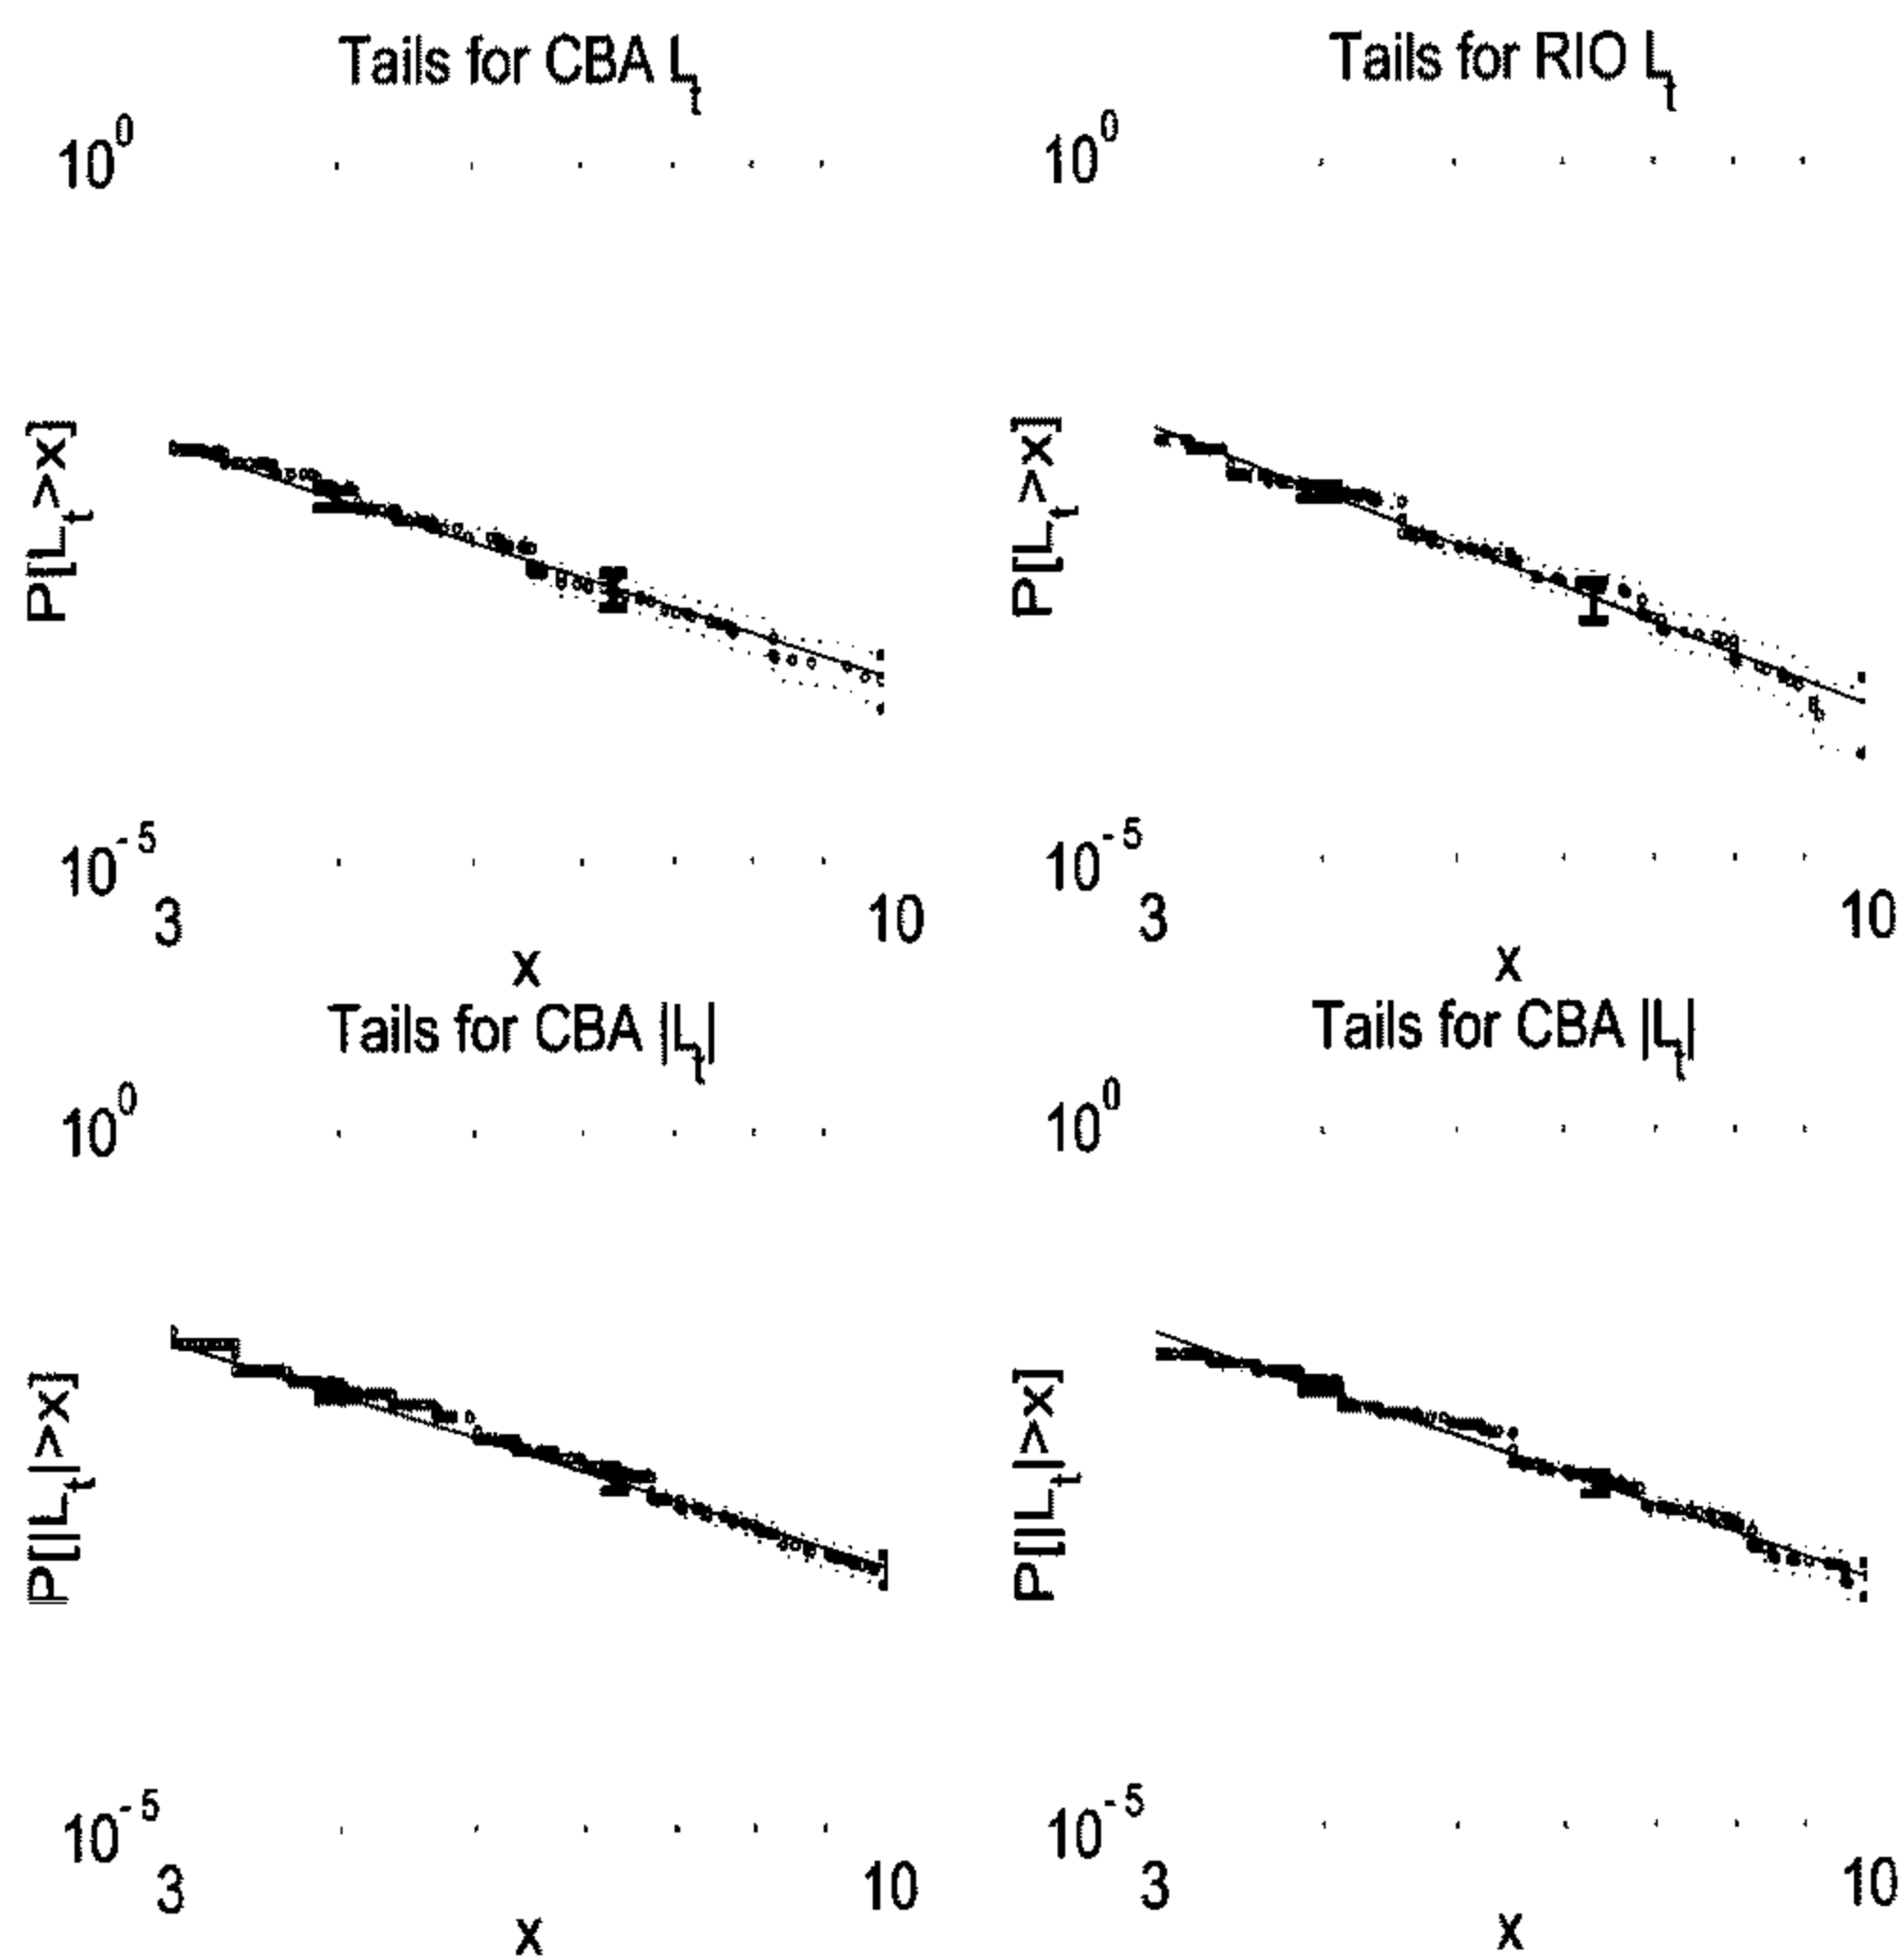


Figure 3.24: The above plots show the tails for CBA and RIO $L_{t,\tau}$ and $|L_{t,\tau}|$, along with estimates for the power law tail behaviour. The dotted lines represent 99% confidence intervals for the distribution tail estimate.

et.al. (2003) have investigated financial data for power law behaviour. The great interest in power law behaviour of financial time series can be attributed to the research of statistical physicists who have long associated power laws with many natural phenomena. Turbulence, diode fluctuation, astrophysics, atmospheric and geophysical systems are some of the traditional fields from where much of this work has been drawn. Power laws are often associated with scaling behaviour, as a power law distribution lacks a characteristic scale.

Studying the empirical distributions of the ASX data we examine the tails of the distribution function for the log-return, absolute log-return, volume and transaction frequency. If the time series display power law behaviour,

$$\mathbf{P}[L_t > x] \sim x^{-\alpha_{L_t,\tau}} \quad (3.13)$$

$$\mathbf{P}[|L_{t,\tau}| > x] \sim x^{-\alpha_{|L_{t,\tau}|}} \quad (3.14)$$

$$\mathbf{P}[V_{t,\tau} > x] \sim x^{-\alpha_{V_{t,\tau}}} \quad (3.15)$$

$$\mathbf{P}[N_{t,\tau} > x] \sim x^{-\alpha_{N_{t,\tau}}}, \quad (3.16)$$

we can estimate the power law indices, $\alpha_{L_t,\tau}$, $\alpha_{|L_{t,\tau}|}$, $\alpha_{V_{t,\tau}}$, and $\alpha_{N_{t,\tau}}$ by plotting the distribution function on a log-log scale and performing a linear fit.

In Fig. 3.24 we show the left and right tails of the intraday returns $L_{t,\tau}$ for two typical stocks in the data set, Commonwealth Bank (CBA) and Rio Tinto (RIO), sampled with $\tau = 10$ minutes. The tails are shown together with 99% confidence intervals. Estimates for the power law index for our data set yields values around $\alpha_{L_t} \approx 3.6$. This value is well outside that of the Lévy-stable distributions. Table B.10 lists the estimated power law index values for 10 stocks in the data set. Least squares fitting was performed over different tail ranges (measured in standard deviations). The best power law behaviour, measured using Pearson's correlation coefficient, was found by fitting the distribution tails across the range of $3 < \sigma < 15$ standard deviations as shown in Table B.10. Increasing the sampling time, τ , we find that the power law index remains unchanged as indicated by Fig. 3.25. This result shows the scale invariance of the ASX data and differs from the findings of studies such as Gorski et.al.(2002) who found that tail exponents do change with τ in the DAX index log-return.




Figure 3.25: This plot shows the distribution tails for CBA for different time interval lengths τ . We can see that there is little change in the distribution properties over sampling interval.

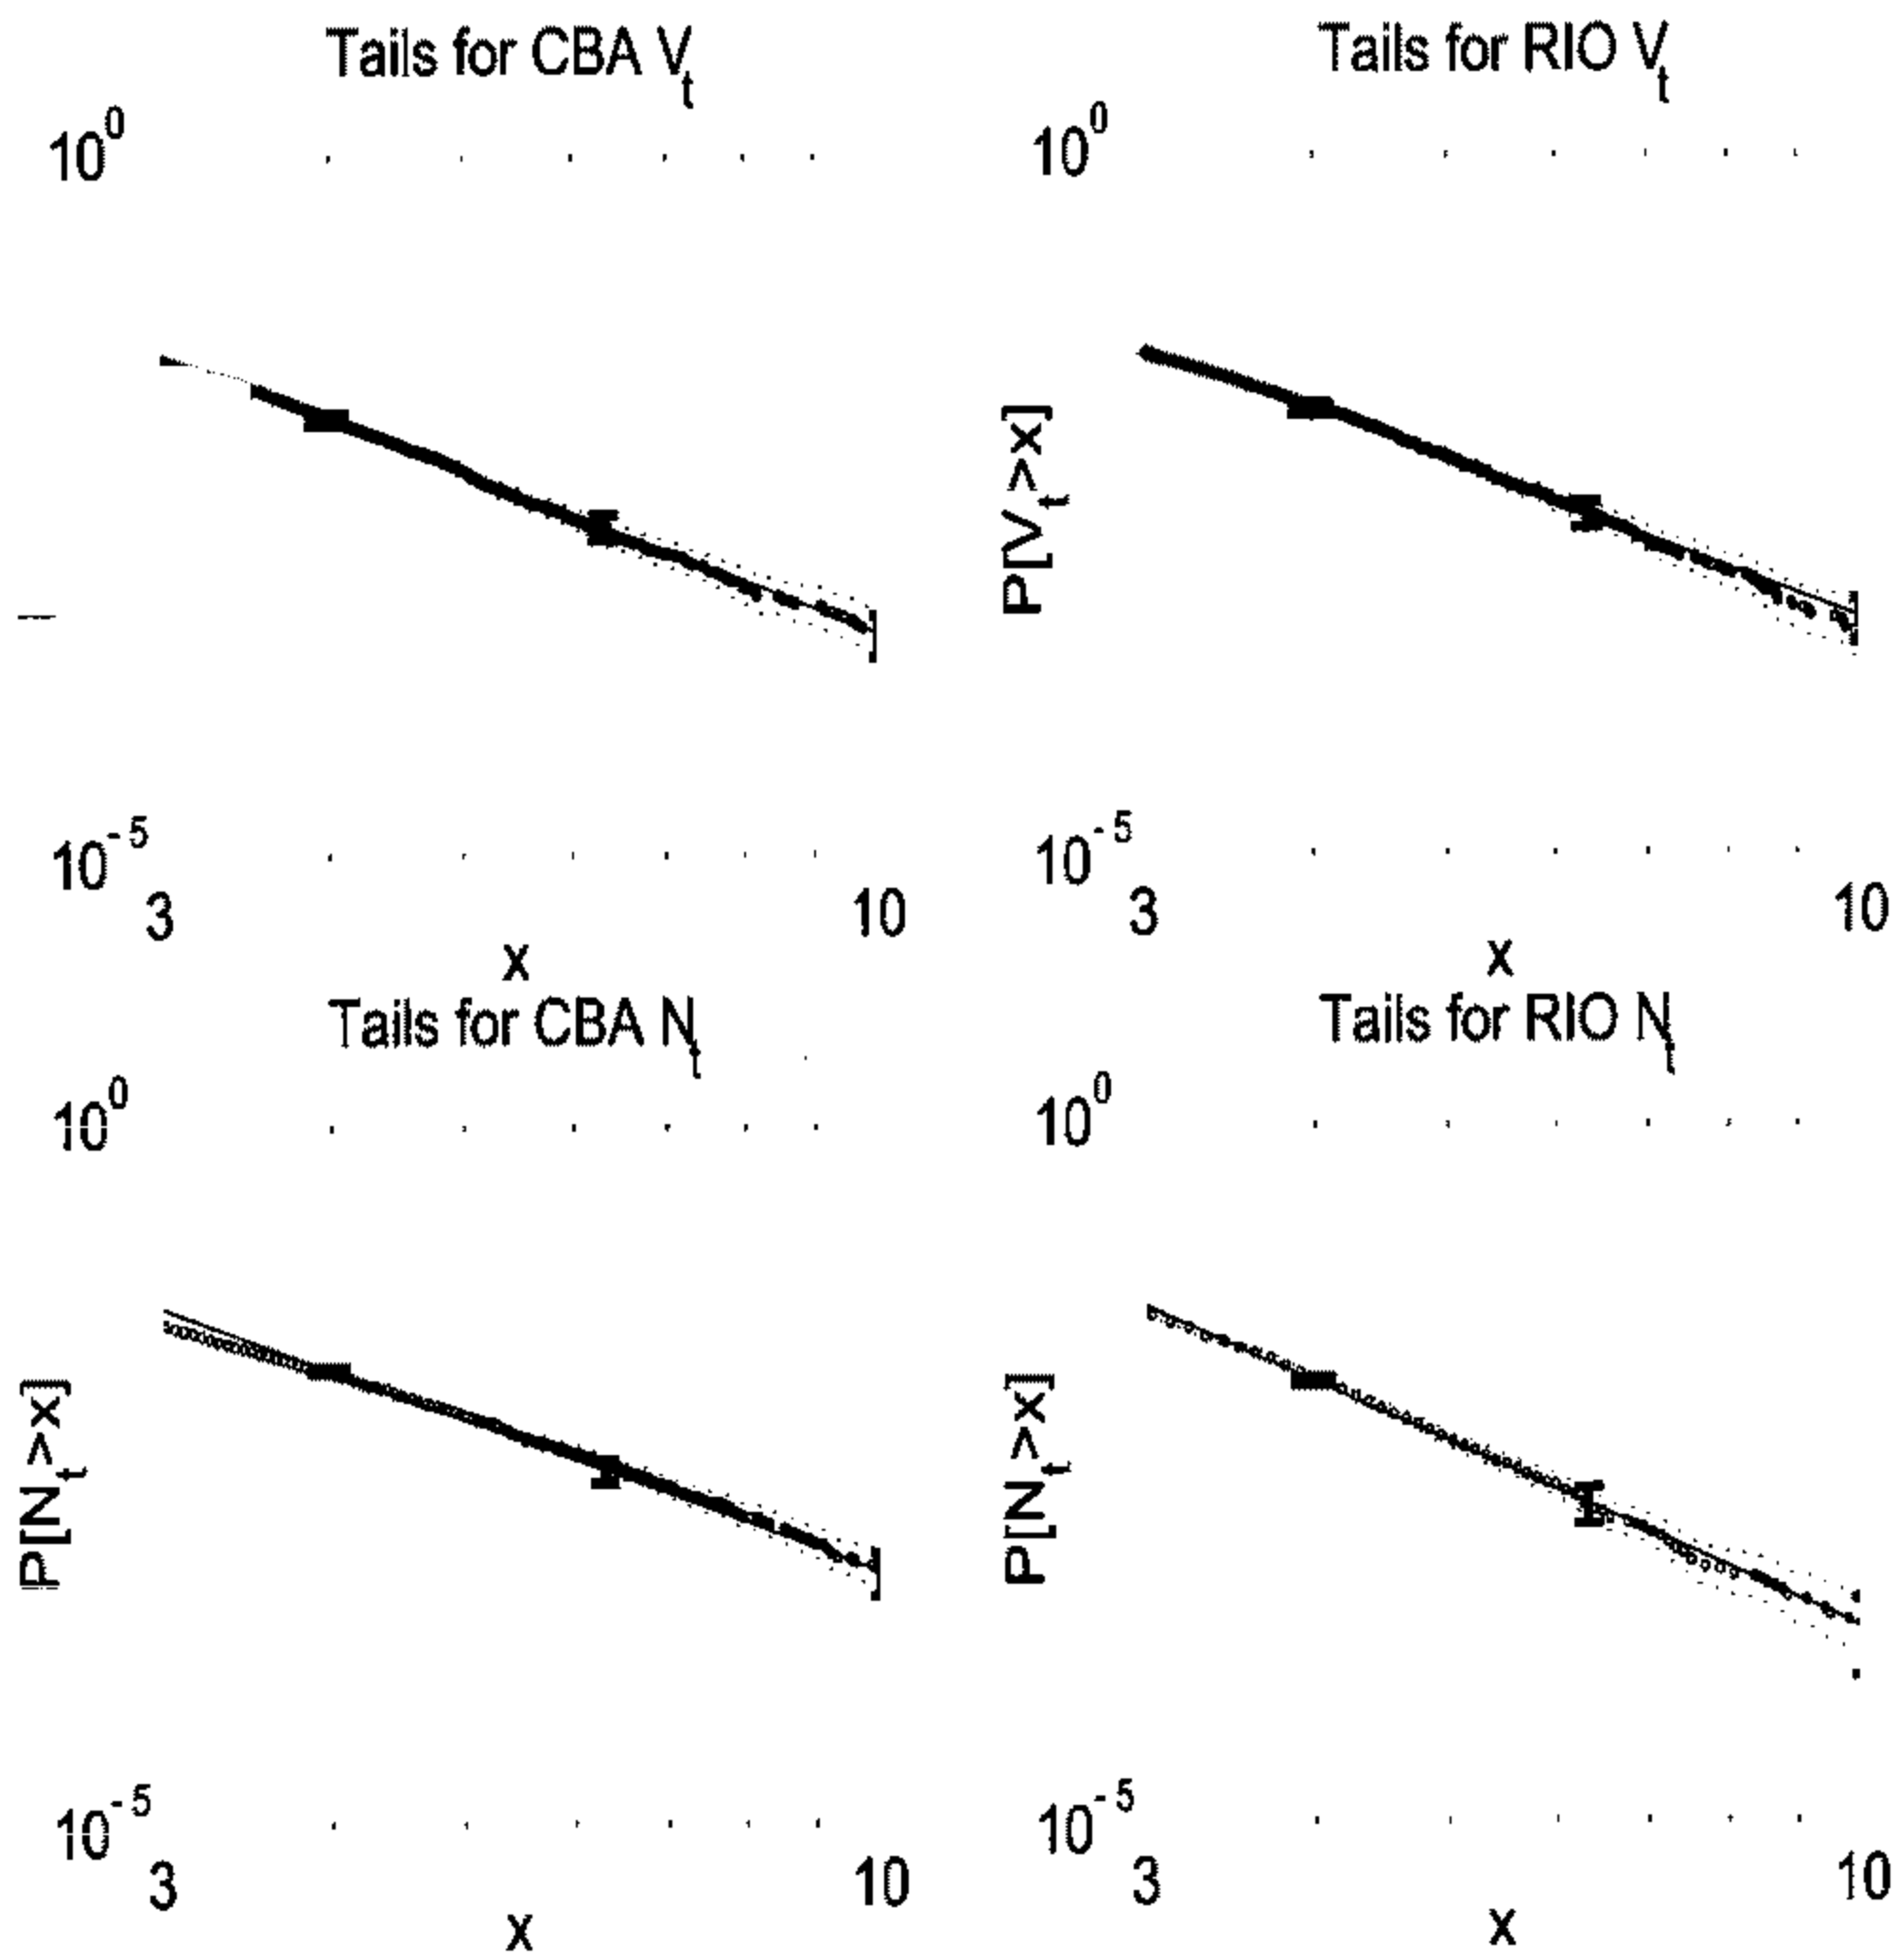


Figure 3.26: The above plots show the tails for $V_{t,T}$, $N_{t,T}$ for both CBA and RIO. These series are proxies for the market volatility and show strong power law type behaviour.

researchers is the behaviour of $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$: the volatility, and hence measures of risk for the τ we find power law index values of $\alpha_{|L_{t,\tau}|} \approx 3.6$, as shown in Fig. 3.24, Fig. 3.26 and Table B.11. This for the ASX data differs from values observed in (2003a), Gabaix et.al. (2003b) and Gopikrishnan et.al. (2003c) where the traded volume was found to possess a tail index reported for the S&P 500. Also, the relative values of L_t , $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ were found to differ from each other. As a result, a recent explanation of financial power law behaviour by Gabaix et.al. (2003a), based on the behaviour of heteroskedasticity, is contradicted by our analysis. That study hypothesises that $|L_{t,\tau}| \sim k\sqrt{V_{t,\tau}}$, with k a constant. This is clearly not in agreement with the behaviour detected in the ASX data set.

3.3.4 The Hill estimator

The Hill estimator was introduced by Hill (1975) in order to provide information on distribution tail behaviour of the form,

$$F(x) \sim 1 - C x^{-\alpha} \tag{3.17}$$

for large x . The Hill estimate is a popular method for determining the power law index, α , for data obeying a Pareto type distribution.

The Hill estimator is constructed as follows. Given n observations of an event, X_1, X_2, \dots, X_n , we sort the observations in order of decreasing magnitude,

$$X_{(1)} \geq \dots \geq X_{(n)}. \tag{3.18}$$

Here $X_{(i)}$ is known as the i^{th} order statistic. The Hill estimator is defined as,

$$H_{n,k} = \frac{1}{k} \sum_{i=1}^k \log \frac{X_{(i)}}{X_{(k+1)}} \tag{3.19}$$

for $k = 1, \dots, n - 1$. The index estimate is obtained as,

$$\alpha_{\text{Hill}} = \frac{1}{H_{n,k}}. \tag{3.20}$$




Figure 3.27: Hill estimates for random samples drawn from the Pareto, Student- t , Gaussian and Lévy distributions. In the case of the Pareto and Student- t , the Hill estimate gives an accurate measure of tail index. For the Gaussian the estimate decays to zero since there is no power law behaviour. For the Lévy distribution the Hill estimate is not able to return an accurate estimate.

The standard error of the estimator is given by,

$$\text{STD}(\alpha_{\text{Hill}}) = \frac{k\alpha_{\text{Hill}}}{(k-1)(k-2)^{1/2}}. \quad (3.21)$$

The Hill estimator has been studied by Hall (1982) and Dekkers and De Haan (1993) and is a popular method to determine the tail index of a distribution with Pareto type tails. For $0 < \alpha < 2$ the Lévy-stable distribution is of this type and the Hill estimator has been applied to it in studies such as Rachev and Mittnik (2000).

Fig. 3.27 shows the results of applying the Hill estimator to several series of random variables, each drawn from a different distribution. The Hill estimates obtained for the Pareto and Student- t distributions both show good agreement with the actual power law index. For the Gaussian case the Hill estimate decreases slowly to zero as there is no power law behaviour. However, for the stable distribution, the estimator does not give a clear result as to the behaviour of the tails. This fact has been examined by Rachev and Mittnik (2000) who proposed using Pickand's estimator as an alternative and serves to illustrate the danger of using a single test to determine the behaviour of the data being examined.

Fig. 3.28 and Fig. 3.29 are plots of the Hill estimate for the tail index for the Australian equity data. In each of the plots we find the Hill estimate seems to approach a constant value, indicating power law behaviour. It can be seen from these plots that the estimated tail index is approximately the same as that found by least squares fits in the previous section. Also note that for $L_{t,\tau}$ and $|L_{t,\tau}|$ the Hill estimate displays a quite jagged line. This is due to the discrete nature of the price data for small returns but seems to have negligible effect on the quality of the estimate. To estimate the power law exponent from the Hill plot we fit a zero order polynomial starting from the 200th order statistic. Table B.12 contains the Hill estimates for the 10 equities examined in the previous section. In each case the Hill estimate confirms the tail index found through empirical estimation of Section 3.3.3. The Hill estimates in Fig. 3.28 do not display the same type of divergent behaviour as that of the Lévy distribution in Fig. 3.27. The tail index estimates found here together with those of the previous section provide a strong argument against stable Paretian hypothesis in favour of finite variance models with power law tails.

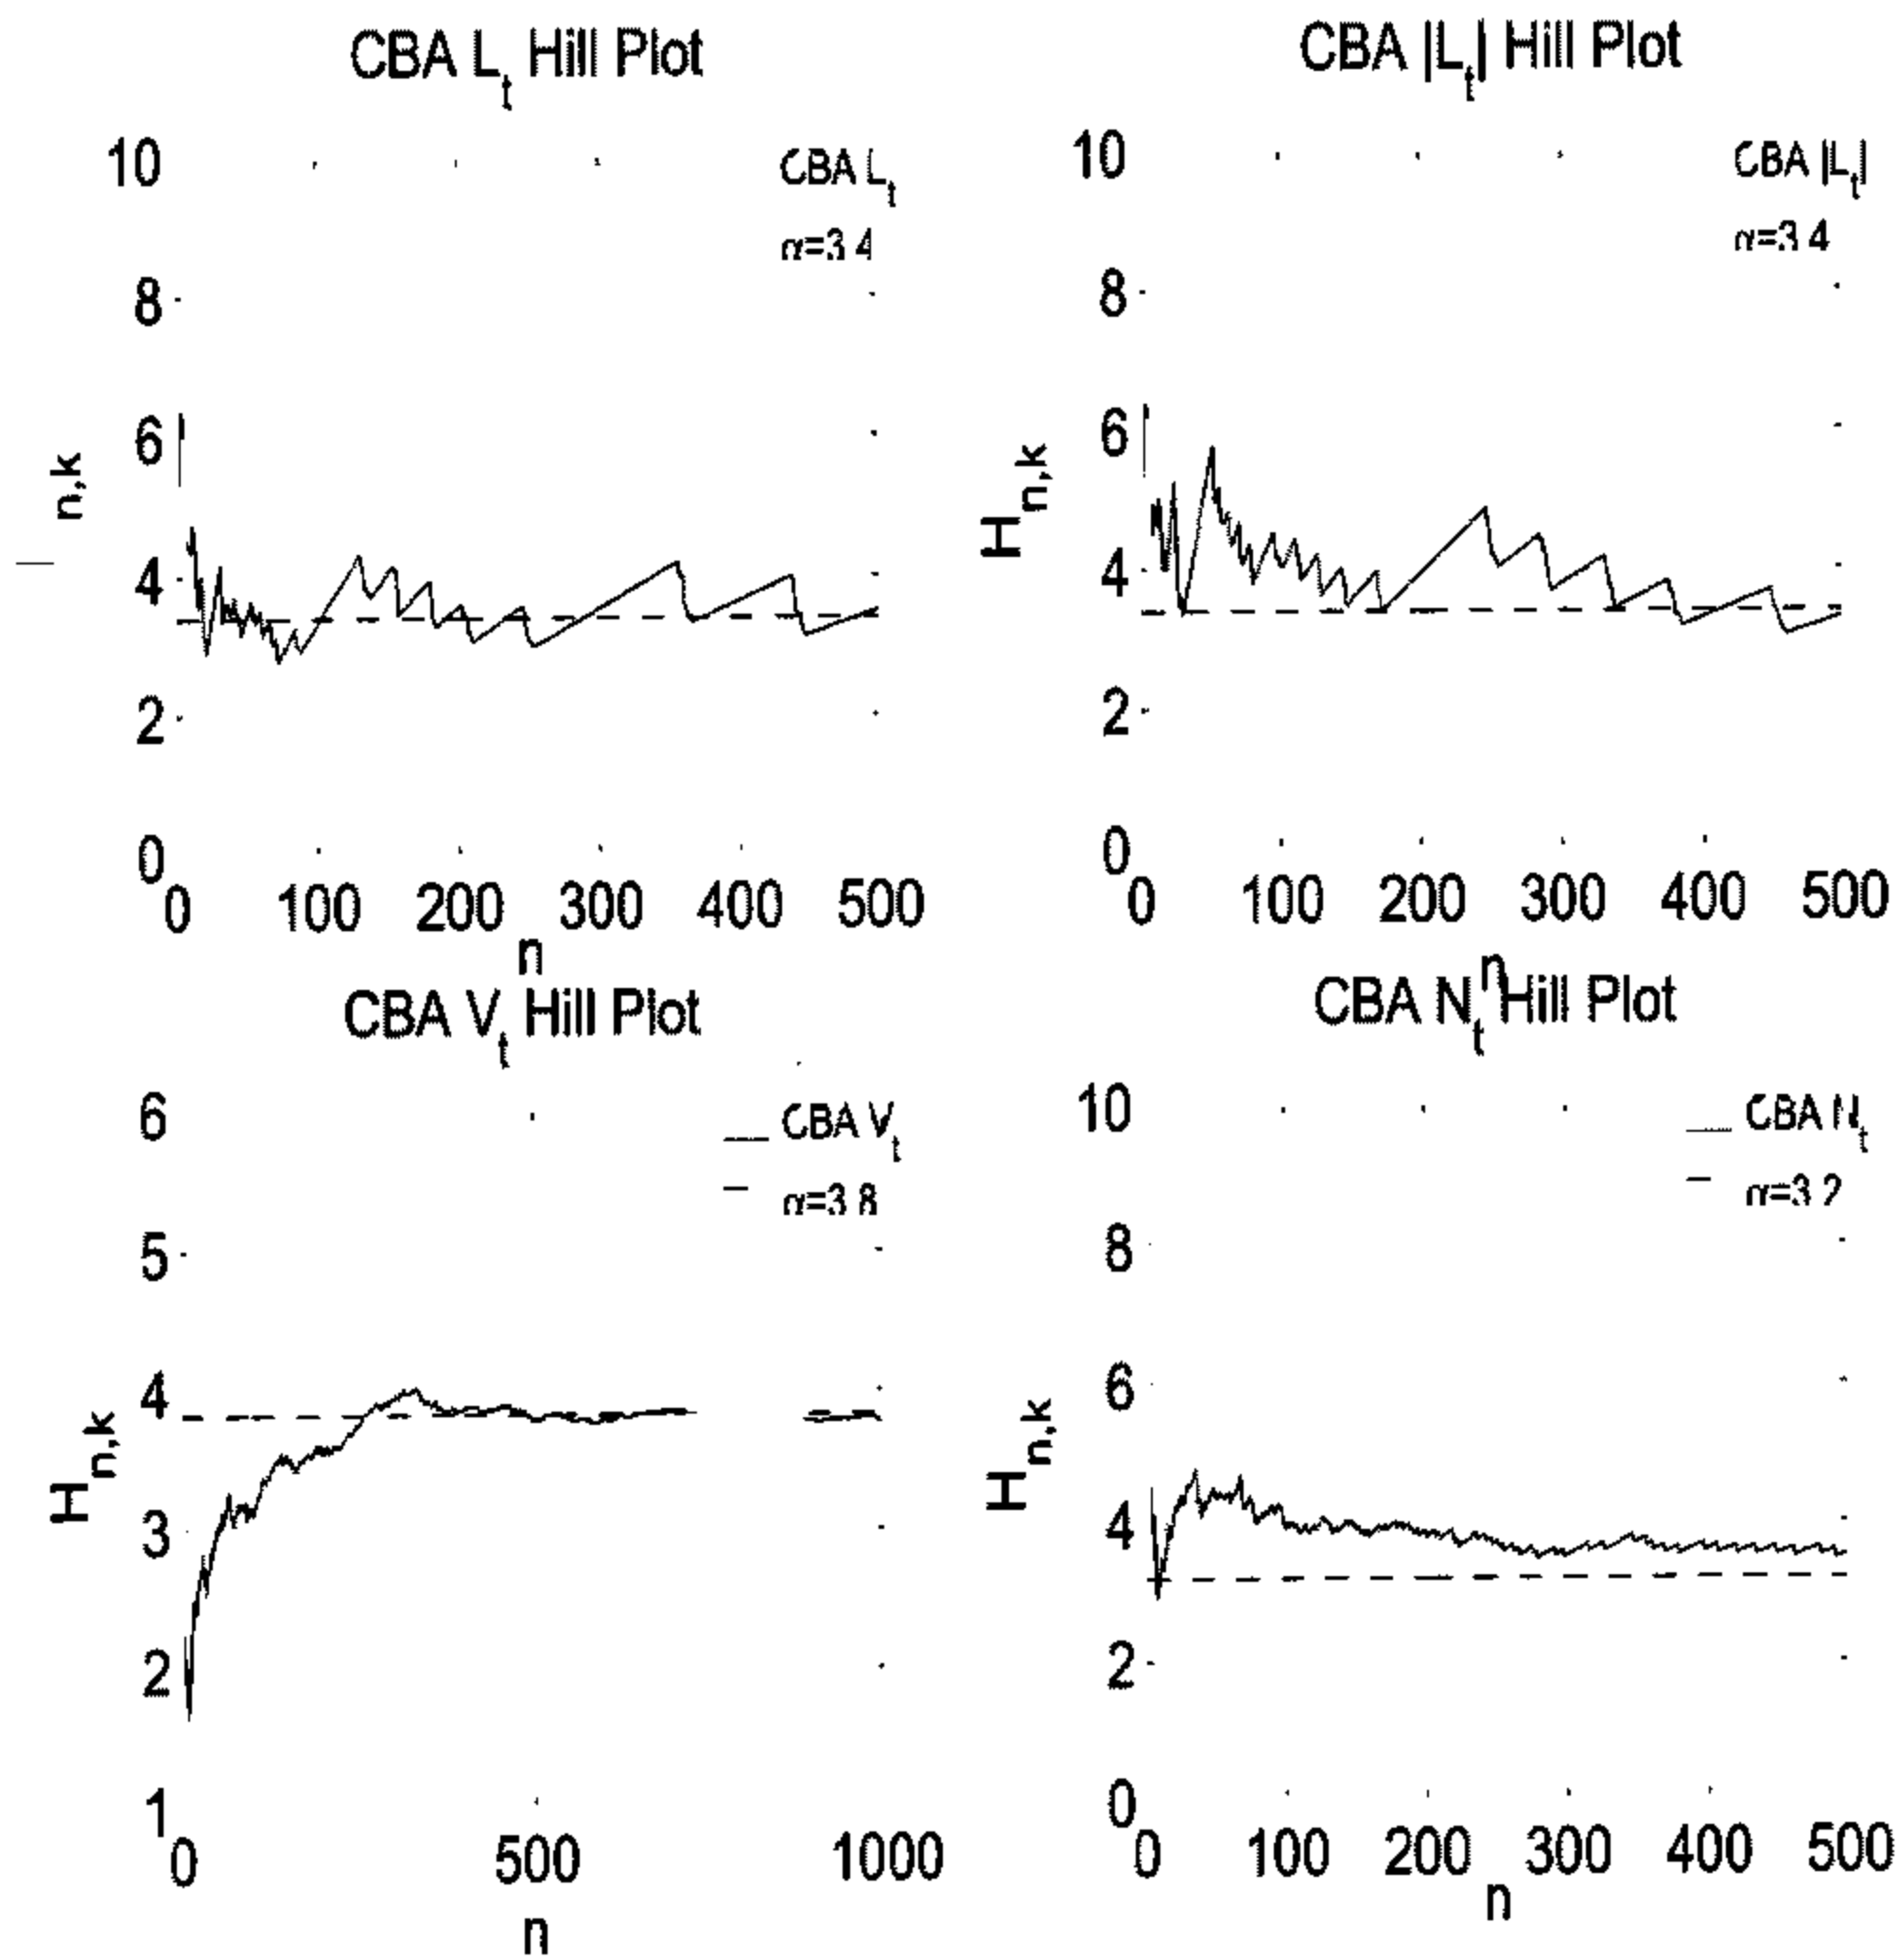


Figure 3.28: Hill estimates for CBA distribution tails. The dashed line represents the empirically estimated indices for $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$.

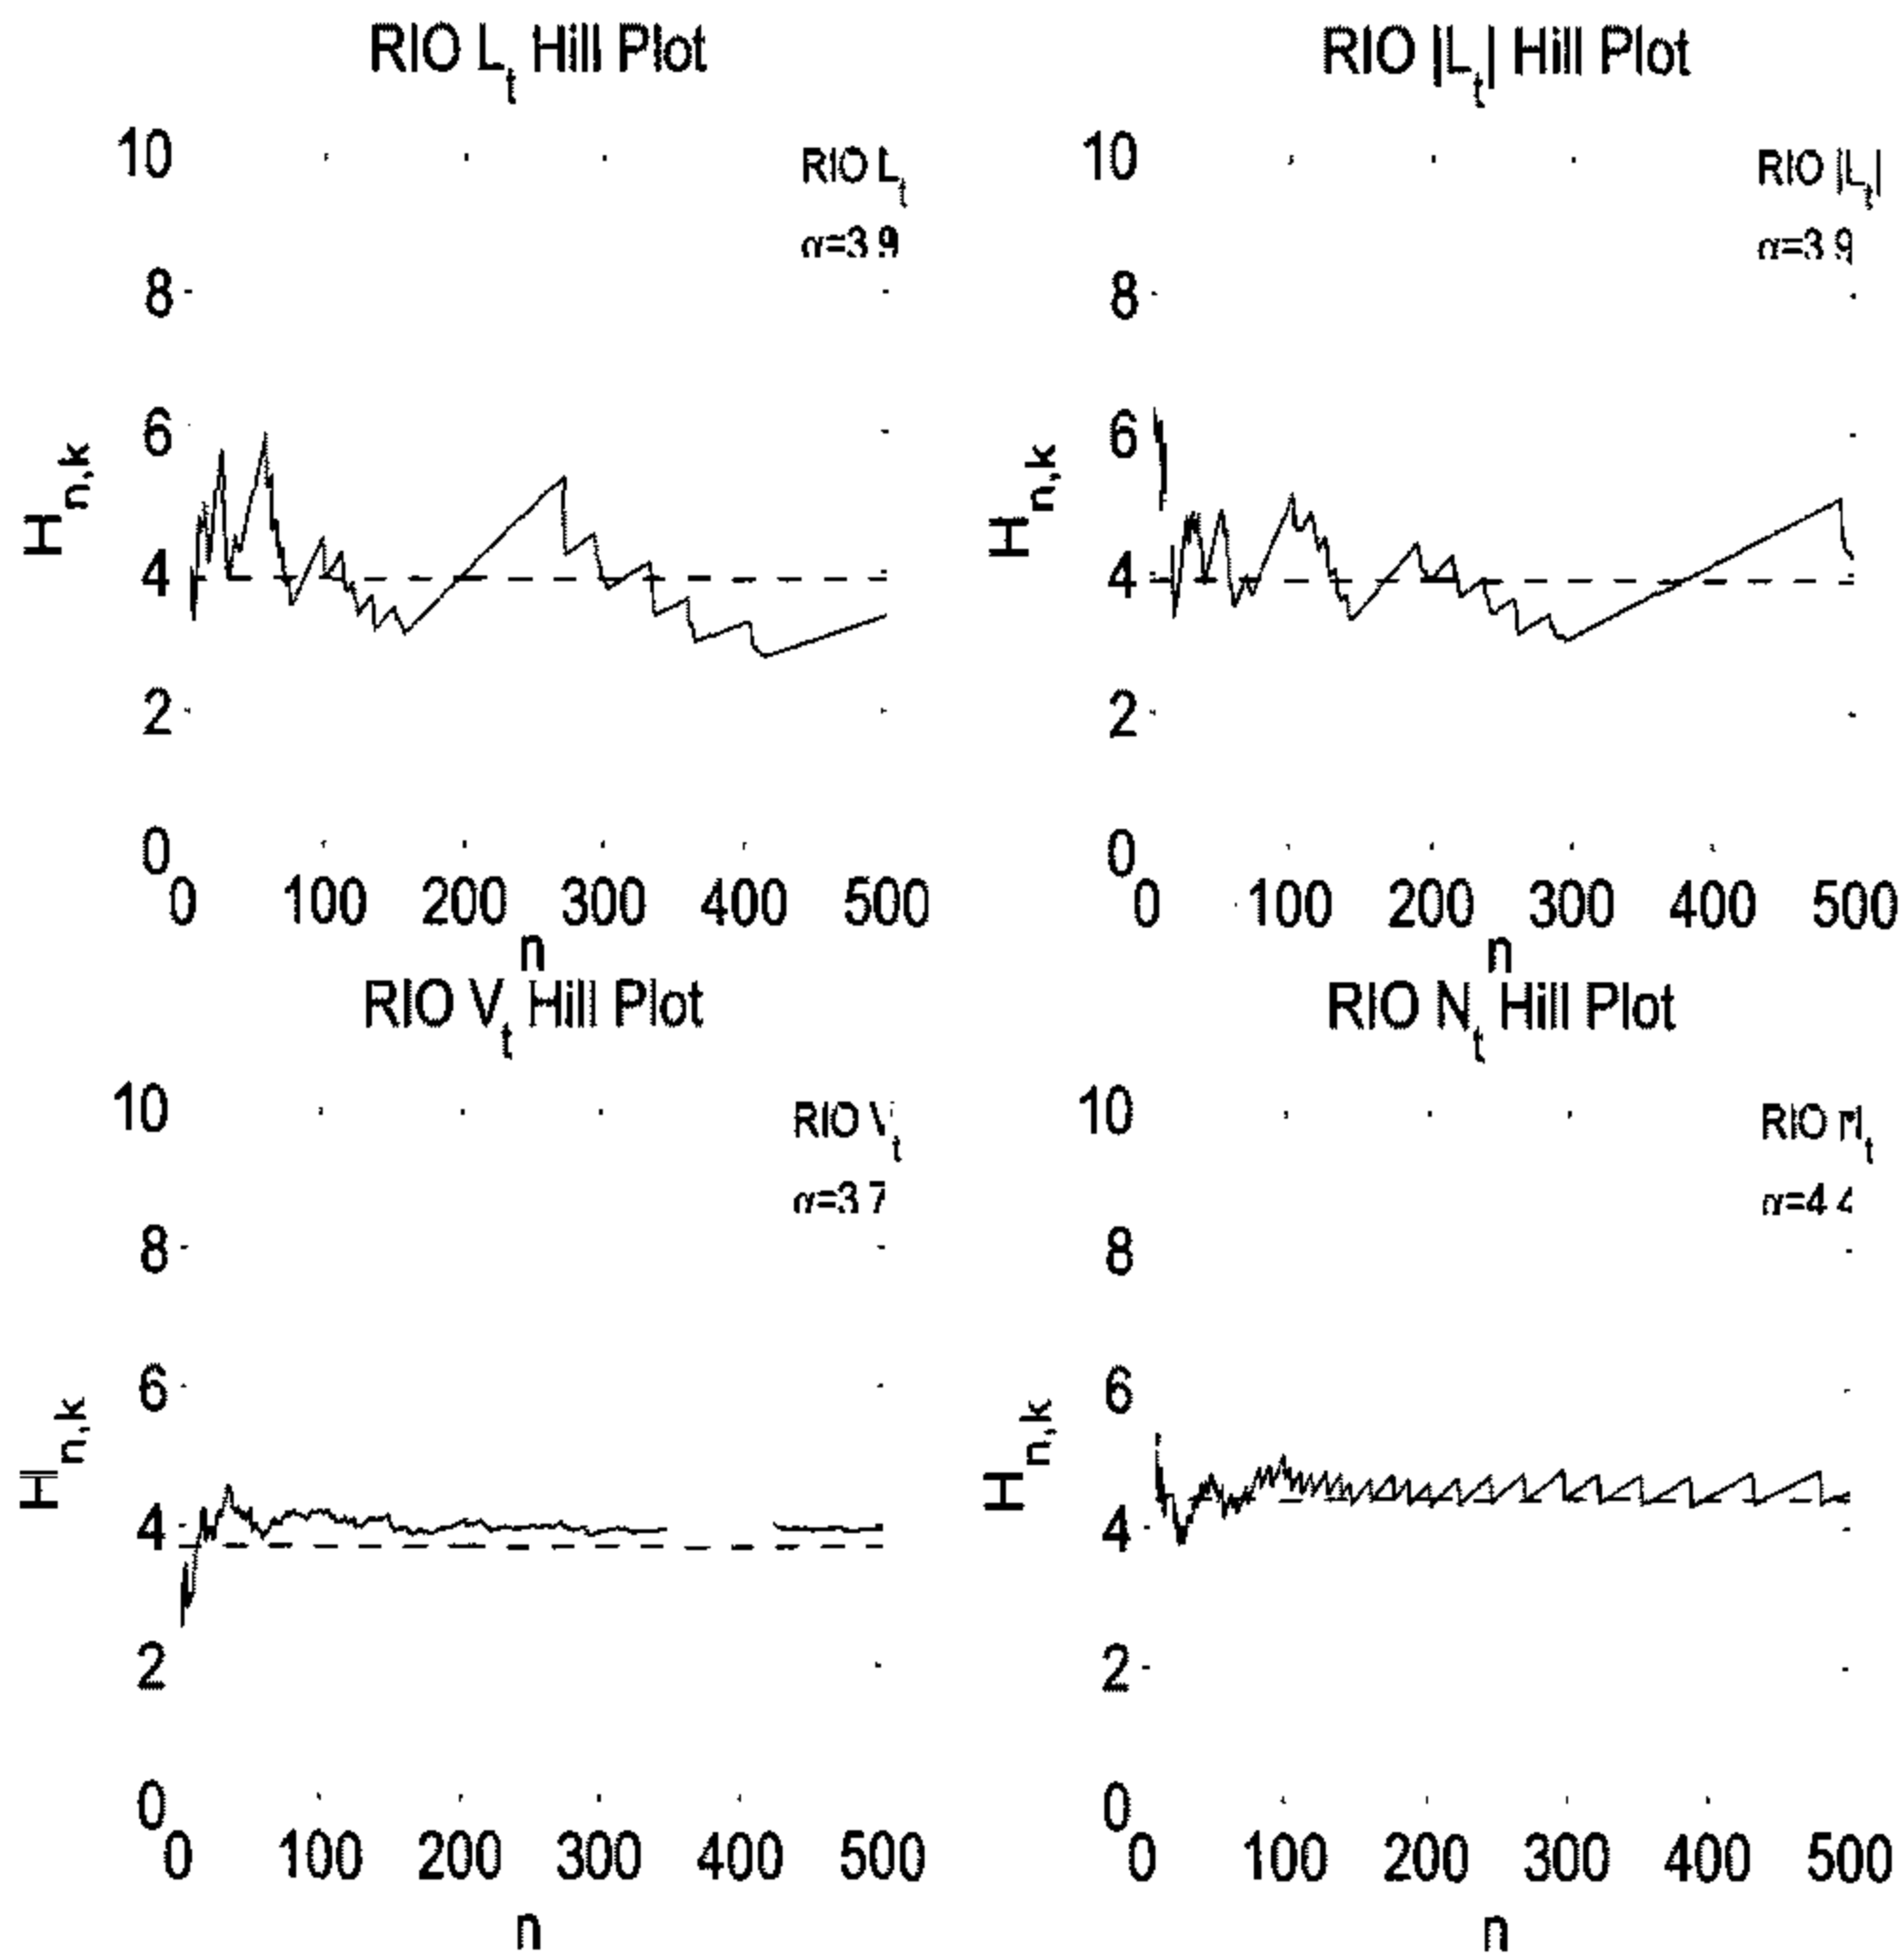


Figure 3.29: Hill estimates for RIO distribution tails. The dashed line represents the empirically estimated indices for $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$.

3.4 Dependency properties

In the previous section we saw that the distribution tails for Australian stock data correspond to a power law with index, $\alpha > 2$. This index value lies clearly outside the Lévy-stable regime, yet the distribution tails seem to be stable over a significant time scale (as shown in Fig. 3.25). For this to occur means that one of the conditions of the central limit theorem is being violated or else the distribution would converge rapidly to Gaussian.

The assumption of independent returns under GBM and other models was originally proposed from an economic point of view. Under the principle of no-arbitrage, any correlation in the price movement will be exploited by traders, whose actions will have the effect of removing the correlation. As we have shown in Chapter 2 and in Section 3.2, although stock returns appear to be independent, they exhibit higher order correlations. In addition, we also find such correlations in measures of the volatility such as the volume and transaction frequency. Fitting the GARCH(1,1) model to high frequency data, we found (see Fig. 3.20) that the exponential correlation structure of the model decays too fast to properly describe the dependence structure.

In this section we will examine the dependence structure of the ASX data. In particular we will investigate the phenomenon of long range dependence (LRD) using the methods described in Beran (1996). The importance of LRD in economics was first recognised by Granger (1966). In finance the study of LRD has been a topic of research for a long time, yet it still remains a controversial issue as to whether asset models should incorporate this behaviour. Here we set out to show that ASX equities contain LRD features.

A process is said to be long range dependent, or possess long memory, if it has autocorrelation function of the form,

$$\rho(k) \sim c_\rho |k|^{-\beta} \quad (3.22)$$

with $0 < \beta < 1$, c_ρ is a constant, and k is the lag. A process with such a correlation structure indicates that the dependence between far apart events diminishes very slowly with increasing lag. In fact, looking at the area under the autocorrelation

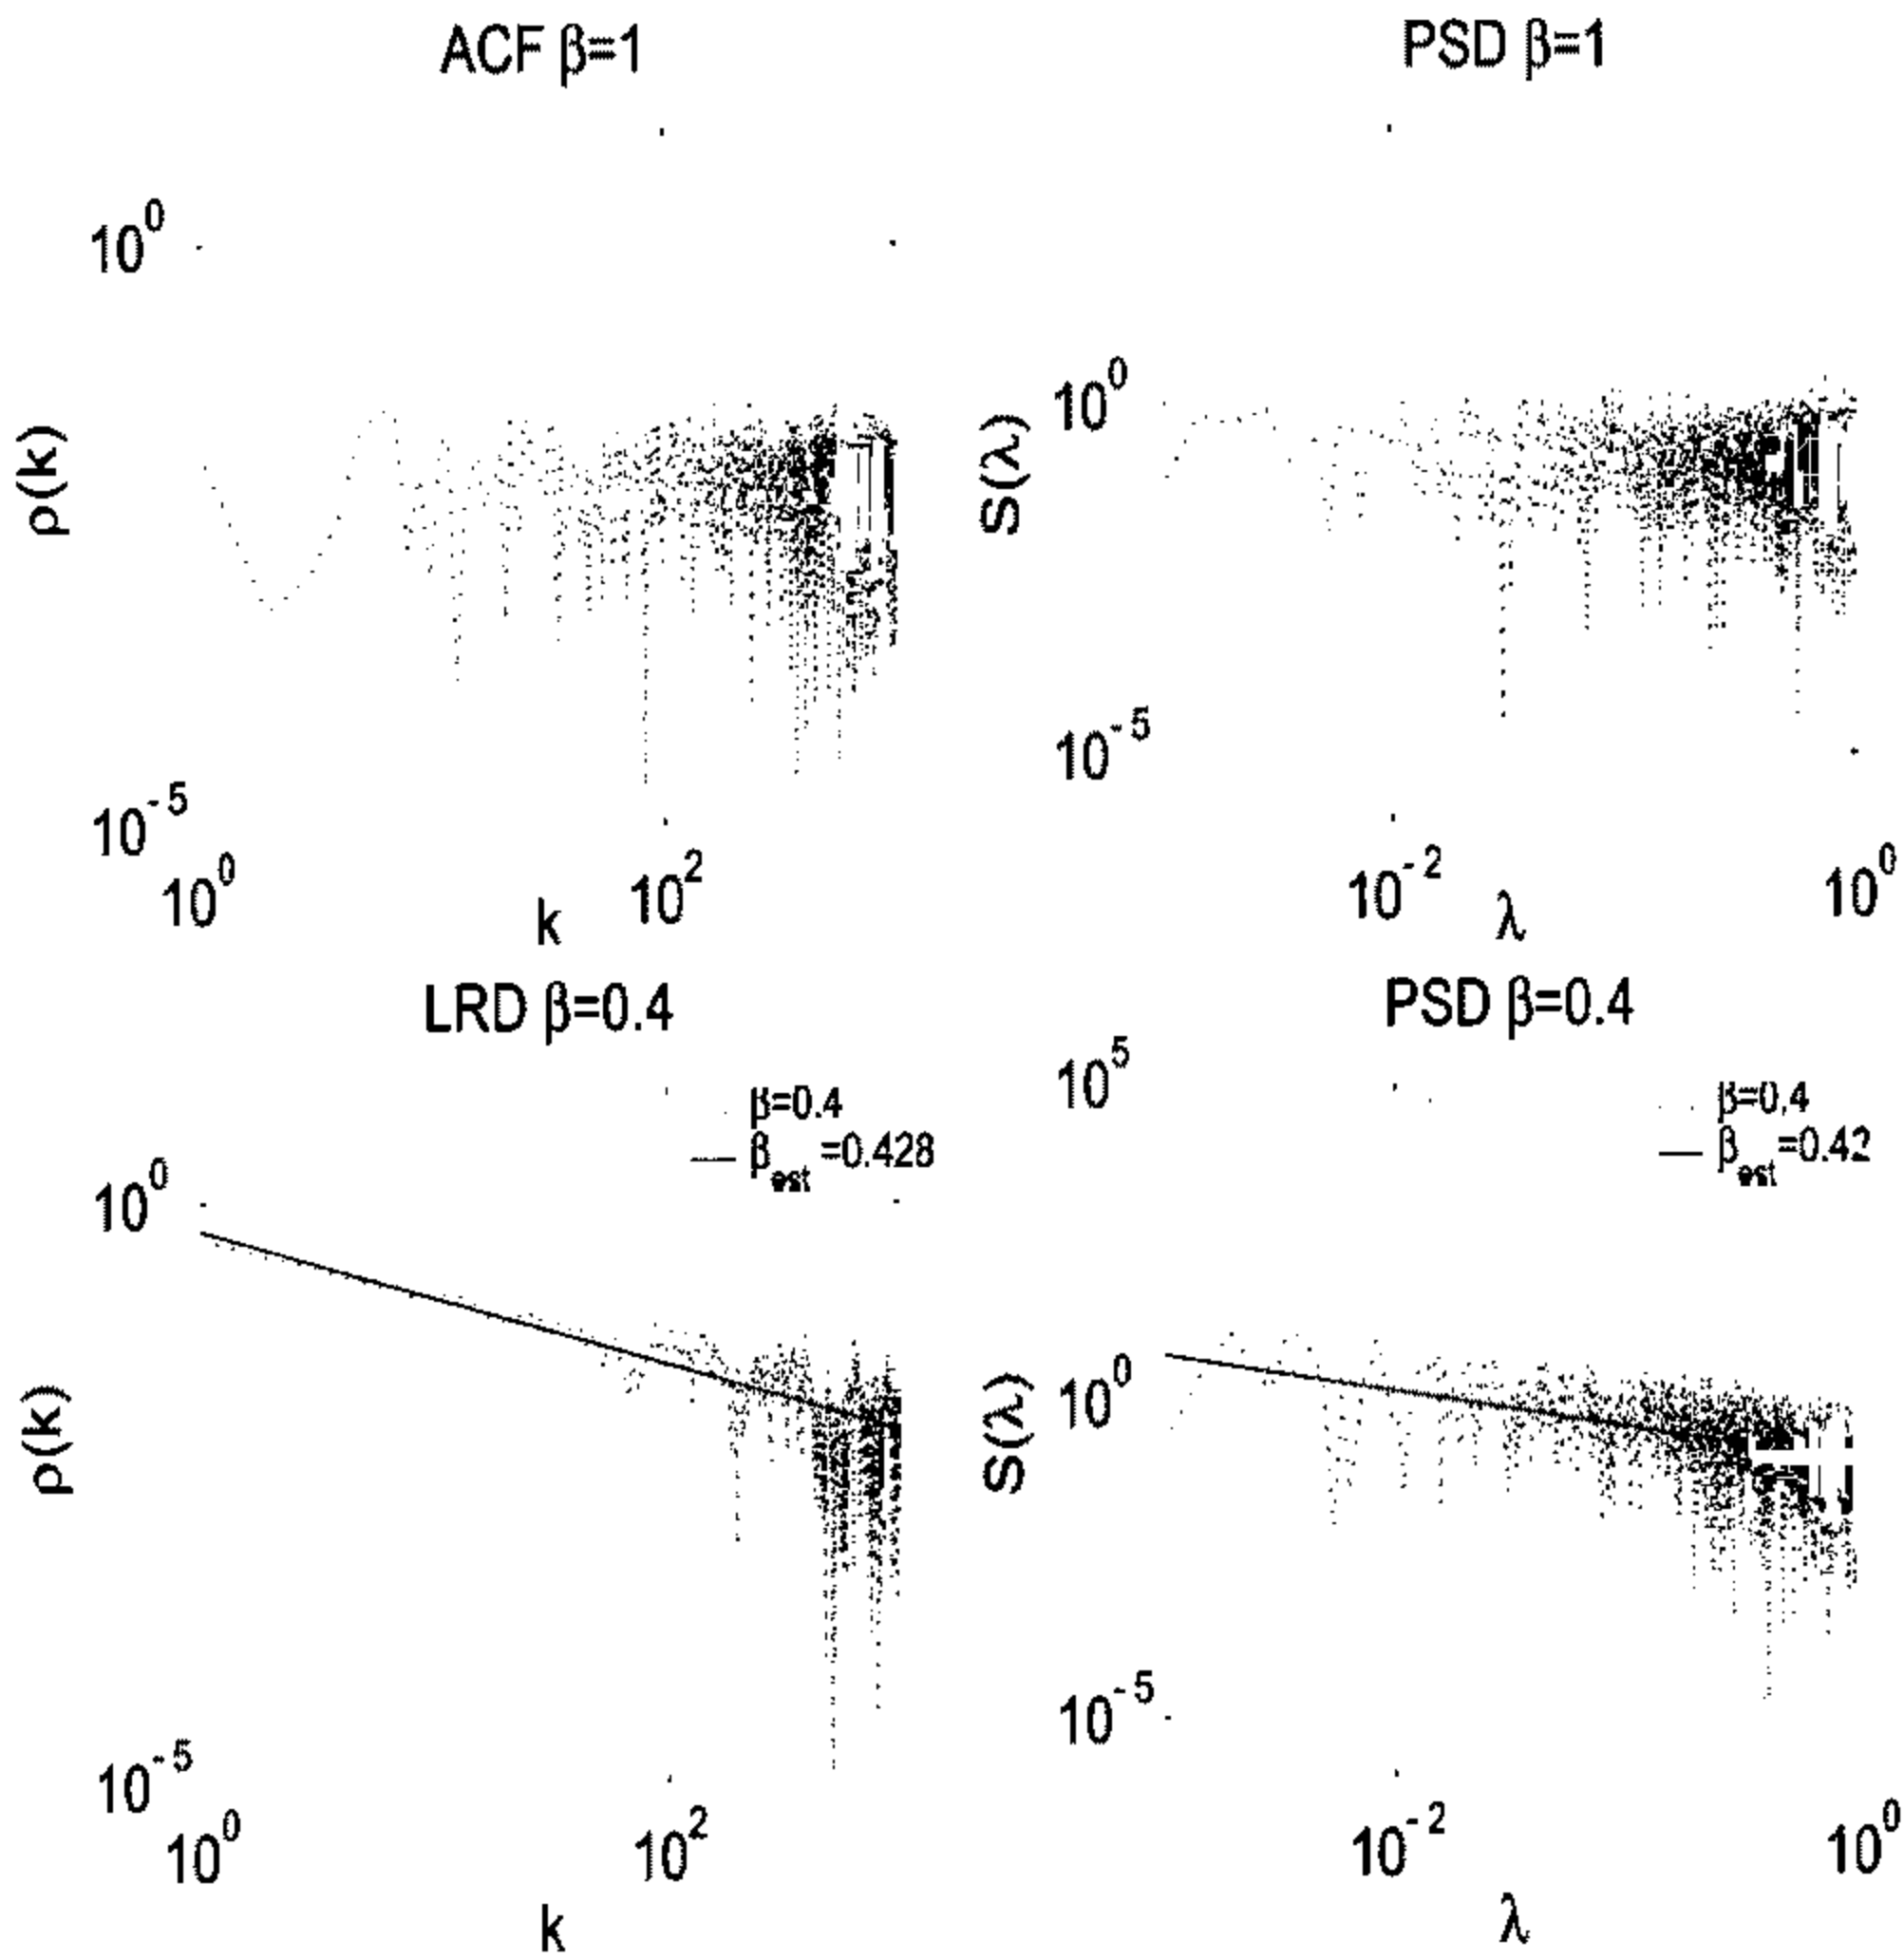


Figure 3.30: The ACF and PSD estimates for uncorrelated and LRD random variables. For the uncorrelated case the ACF is that of white noise while the PSD returns $\beta \sim 1$. For LRD, the ACF and PSD provide estimates of the parameter β .

function,

$$\int_0^{\infty} \rho(k) dk = \int_0^{\infty} c_{\rho} |k|^{-\beta} dk \rightarrow \infty \quad (3.23)$$

for $0 < \beta < 1$. Hence this property is often referred to as infinite memory.

3.4.1 Autocorrelation and periodogram

A suitable test for long memory is obtained by plotting the autocorrelation function on a log-log scale,

$$\log \rho(k) \sim -\beta \log |k| + C \quad (3.24)$$

where C is constant. Thus if long memory is present the plot will show a linear trend with negative gradient.

Alternatively, LRD can be examined in the frequency domain. The spectral density of a process with correlation given by Eq. (3.22) is

$$S(\lambda) = \int_{-\infty}^{\infty} \rho(k) e^{2\pi i k \lambda} dk \quad (3.25)$$

$$\sim c_{\lambda} |\lambda|^{-\eta} \quad (3.26)$$

with $\eta = 1 - \beta$. Thus for $\eta < 1$ such a process is characterised by having a pole at the origin in the frequency domain. Plotting Eq. (3.26) on a log-log scale,

$$\log S(\lambda) \sim (\beta - 1) \log |\lambda| + C, \quad (3.27)$$

we obtain an indicator of the existence of LRD and an estimate of the index β . To illustrate the phenomenon of LRD we have simulated two processes: a white noise process with $\beta = 1$; and a long memory process with $\beta = 0.4$. Fig. 3.30 shows an example of the two LRD estimates applied to the simulated time series. We see that for the white noise process, the ACF and PSD estimates indicate an uncorrelated time series, while for the simulated LRD process we find that the ACF and PSD estimates are able to measure the parameter β .

In Fig. 3.31 is shown the autocorrelation (ACF) and spectral density (PSD) of the ANZ log-return, $L_{t,\tau}$ and absolute log-return, $|L_{t,\tau}|$. A linear least squares fit has been performed in each case to estimate the LRD parameter β . It is seen that




Figure 3.31: ACF and PSD estimates for ANZ $L_{t,\tau}$ and $|L_{t,\tau}|$. In the case of $L_{t,\tau}$ the ACF and PSD indicate an uncorrelated time series. For $|L_{t,\tau}|$ the estimates indicate the presence of long memory.

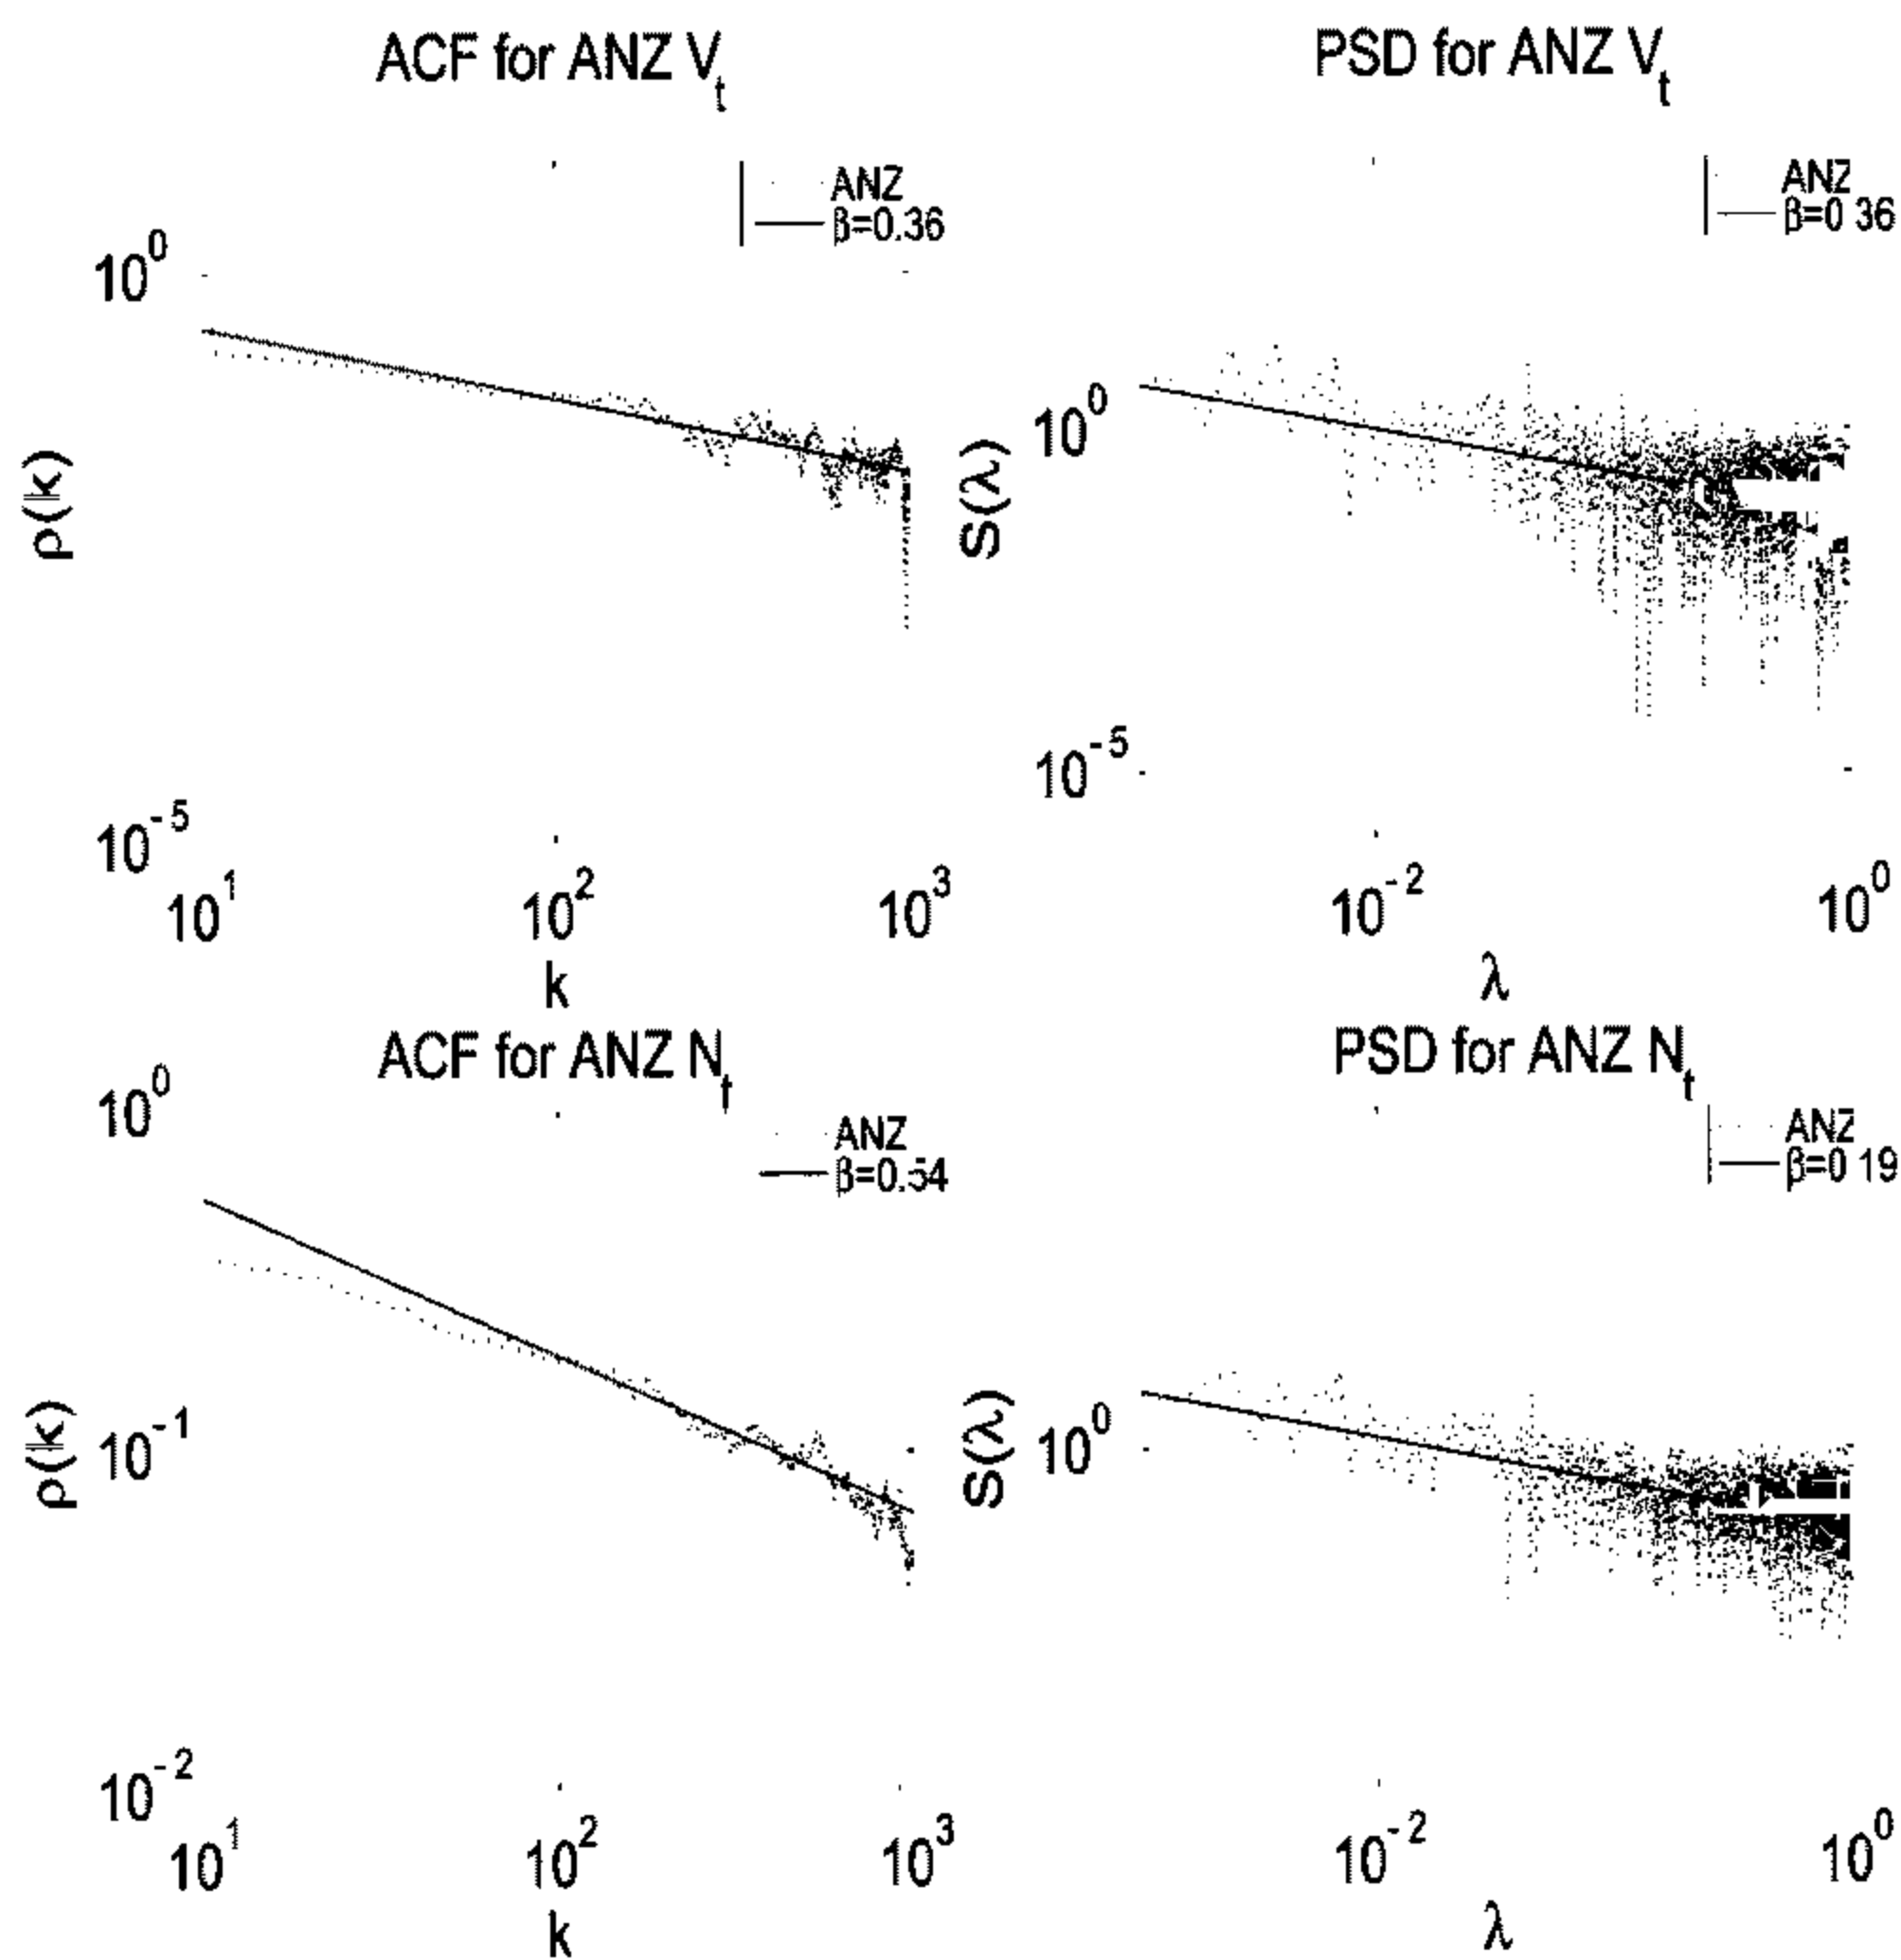


Figure 3.32: ACF and PSD estimates for ANZ $N_{t,\tau}$ and $V_{t,\tau}$. For both series the estimates indicate the presence of long memory.

for both estimation methods, the results indicate uncorrelated random variables for $L_{t,\tau}$. However, the ACF and PSD estimates for $|L_{t,\tau}|$ indicate the presence of long memory. Fig. 3.32 displays the ACF and PSD estimates for ANZ volume, $V_{t,\tau}$, and transaction frequency, $N_{t,\tau}$. Again, a linear least squares fit shows that the correlation structure is well described by a power law with index, β , that falls well inside the range of LRD. Table B.13 and Table B.14 list the results for applying these estimates to the data set.

3.4.2 Variance plots

A process can be tested for a correlation structure such as Eq. (3.22) by examining the variance of the sample mean (Beran (1996)). This method is known as a variance plot.

The variance of the sample mean of a time series X_t can be represented in terms of its autocorrelation,

$$\mathbb{V}(\bar{X}) = n^{-2} \sigma^2 \sum_{i,j=1}^n \rho(i, j) \quad (3.28)$$

where n is the length of X , $\sigma = \mathbb{V}(X)$ and $\rho(i, j)$ is the autocorrelation matrix of X . If $\rho(i, j)$ only depends on $k = |i - j|$ the process is said to be weak-stationary and we may write

$$\mathbb{V}(\bar{X}) = n^{-1} \sigma^2 \left[1 + 2 \sum_{k=1}^{n-1} \left(1 - \frac{k}{n} \right) \rho(k) \right] \quad (3.29)$$

$$= \frac{\sigma^2}{n} + \frac{\sigma^2}{n^\beta} n^{\beta-1} 2 \sum_{k=1}^{n-1} \left(1 - \frac{k}{n} \right) \rho(k). \quad (3.30)$$

For large n with $\rho(k) \sim c_\rho |k|^{-\beta}$ this becomes

$$\mathbb{V}(\bar{X}) \sim \sigma^2 c(\rho) n^{-\beta} \quad \text{with } c > 0. \quad (3.31)$$

As a test for long memory we have plotted the variance of the sample mean for different lengths n in Fig. 3.33 and Fig 3.34 for ANZ Bank (ANZ). From these plots we are able to estimate values for the long memory parameter β for $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$




Figure 3.33: Variance plots for ANZ $L_{t,\tau}$ and $|L_{t,\tau}|$. In the case of $L_{t,\tau}$ the ACF and PSD indicate an uncorrelated time series. For $|L_{t,\tau}|$ the estimates indicate the presence of long memory.




Figure 3.34: Variance plots for ANZ $N_{t,\tau}$ and $V_{t,\tau}$. For both series the estimates indicate the presence of long memory.

and $N_{t,\tau}$ using least squares. In Table B.15 we present the estimates of β for 10 stocks in the data set. For the log-return process, $L_{t,\tau}$, we find the value $\beta \sim 1$, indicating a lack of long memory in this process. However, the β values for $|L_{t,\tau}|$, $V_{t,\tau}$ and $N_{t,\tau}$ provide a strong indication of the presence of long memory in these processes. The values obtained are in agreement with those found via ACF and PSD estimates.

3.4.3 Rescaled range analysis

Rescaled range (R/S) analysis was first proposed by Hurst (1951) to test for the presence of correlations in empirical time series. R/S analysis has been extensively studied by Beran (1996) and Feder (1988) and implemented in many studies over the years such as Mandelbrot and Wallis (1968) and Costa and Vasconcelos (2003).

To calculate the R/S statistic for a series X_1, \dots, X_T , we start by dividing the data into N non overlapping intervals I_n of equal size τ . On each interval we define the adjusted range,

$$R_n = \max_{1 \leq k \leq \tau} \sum_{i=1}^k (X_{n\tau+i} - \bar{X}_n) - \min_{1 \leq k \leq \tau} \sum_{i=1}^k (X_{n\tau+i} - \bar{X}_n) \quad (3.32)$$

where

$$\bar{X}_n = \frac{1}{\tau} \sum_{i=1}^{\tau} X_{n\tau+i} \quad (3.33)$$

and the standard deviation

$$S_n = \sqrt{\frac{1}{\tau} \sum_{i=1}^{\tau} (X_{n\tau+i} - \bar{X}_n)^2}. \quad (3.34)$$

The rescaled range $(R/S)_\tau$ is then defined as the average of the ratio R_n/S_n over all intervals I_n ,

$$(R/S)_\tau = \frac{1}{N} \sum_{n=0}^{N-1} \frac{R_n}{S_n}. \quad (3.35)$$

For any stationary process with short range dependence the R/S statistic, as a function of interval length, should behave asymptotically as $\tau^{1/2}$. However as Hurst

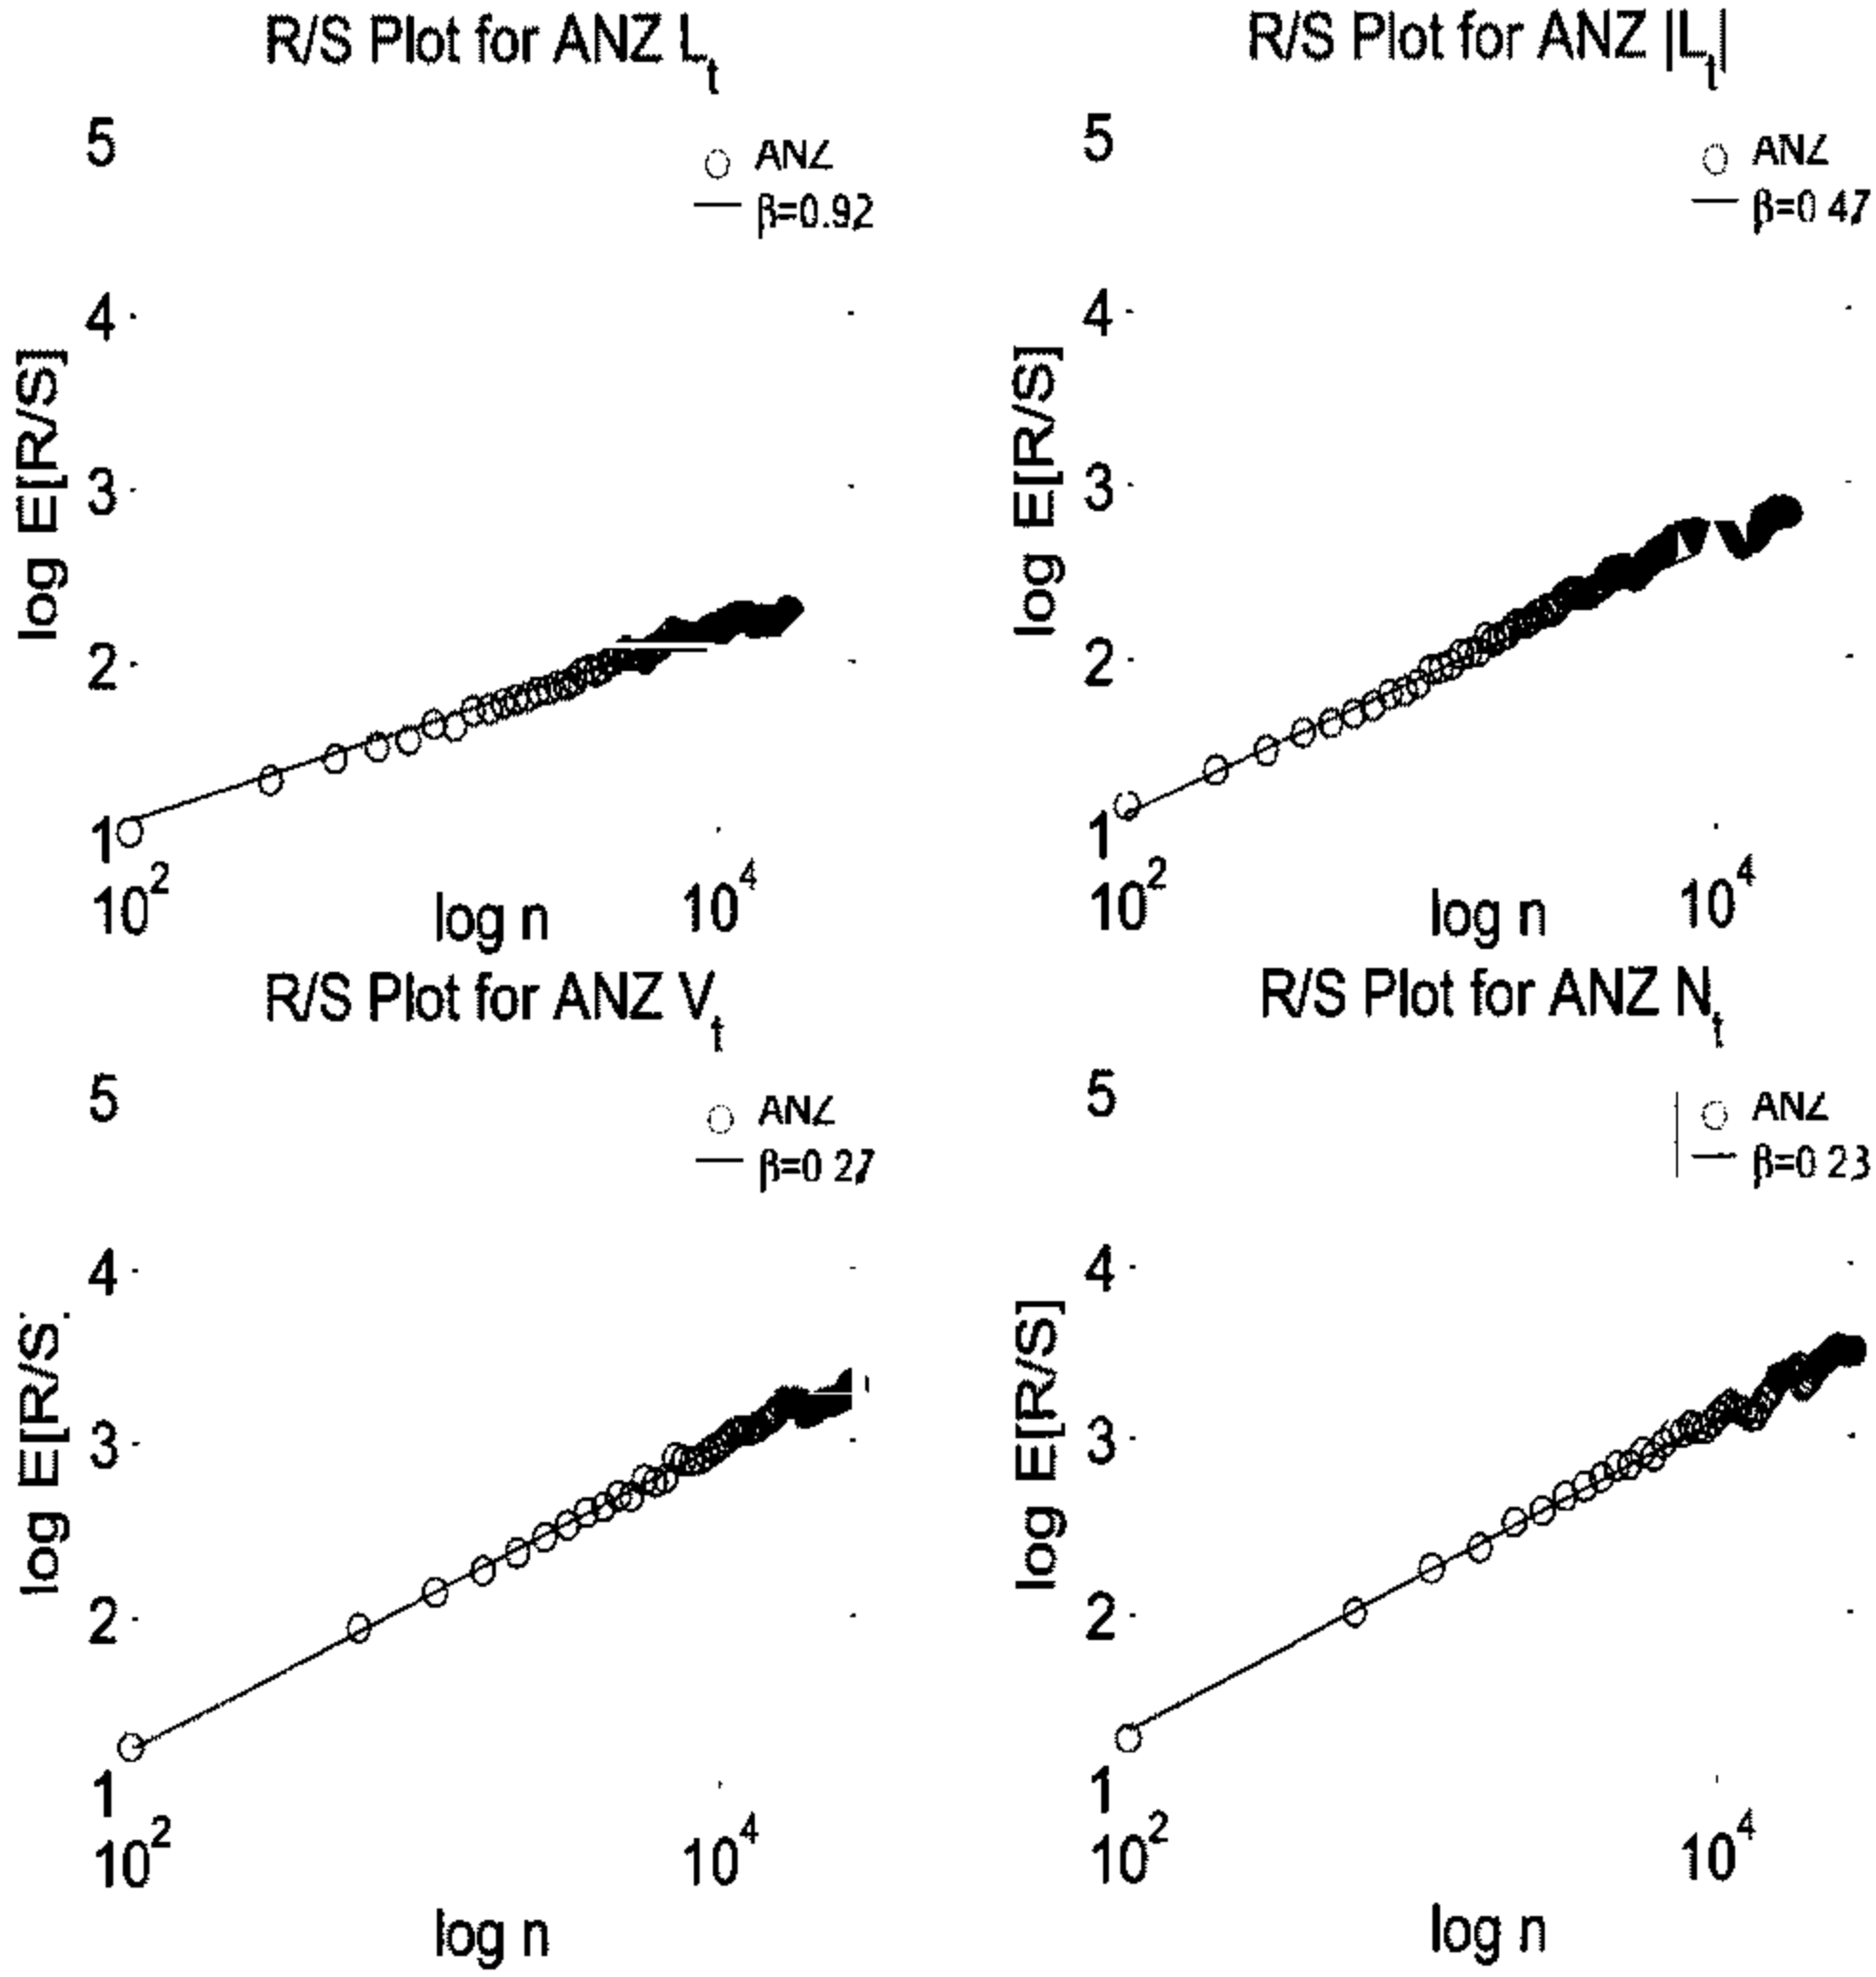


Figure 3.35: R/S analysis for ANZ data.

discovered for an LRD process, plotting the R/S statistic against interval length yields

$$(R/S)_\tau \sim \tau^{1-\beta/2}. \tag{3.36}$$

Thus an estimate of β is gained by linear regression on a log-log scale.

Fig. 3.35 displays the results of performing R/S analysis on ANZ data. Table B.16 contains the results for the studied equity data. Here we see that the R/S analysis indicates that long memory behaviour occurs in the three volatility mea-

tures. Again the results found using this method are in good agreement with those in Tables B.13, B.14 and B.15.

3.4.4 Wavelet test for LRD

The use of wavelets for the detection and estimation of long range dependence was pioneered by Abry, Gonçalves and Flandrin (1995). The wavelet transform conducts a local Fourier analysis by projecting the input signal onto locally oscillating waveforms known as ‘wavelets’. Given a wavelet function $\psi(t)$ with central frequency S_0 , we define the dilations and translations of the wavelet as,

$$\phi_{j,k}(t) = \frac{1}{2^{j/2}}\psi(2^{-j}t - k), \quad (3.37)$$

so that $\phi_{j,k}(t)$ has central frequency $2^{-j}S_0$ and is time shifted by $2^j k$. The discrete wavelet transform (DWT) of a function (or process) $X(t)$ is defined as,

$$d_{j,k} = \int_{\mathbb{R}} X(t)\psi_{j,k}(t)dt. \quad (3.38)$$

The $d_{j,k}$ are known as the wavelet coefficients.

For a stationary process, X_t , with power spectrum $S_X(\lambda)$, the wavelet coefficients satisfy the relationship (Abry et.al. 2000),

$$\mathbb{E}[d_{j,k}^2] = \int S_X(\lambda)2^j|\hat{\psi}(2^j\lambda)|^2d\lambda \quad (3.39)$$

where $\hat{\psi}(\lambda)$ is the Fourier transform of $\psi(t)$. Further, if the power spectrum obeys a power law, $S_X(\lambda) \sim c_f|\lambda|^{\beta-1}$, $0 < \beta < 1$ then Eq. (3.39) becomes,

$$\mathbb{E}[d_{j,k}^2] \sim C(\beta, \psi)2^{j(1-\beta)} \quad (3.40)$$

where $C(\beta, \psi)2^{j(1-\beta)} = \int |\lambda|^{\beta-1}|\hat{\psi}(\lambda)|^2d\lambda$. This behaviour suggests a linear relationship of the form,

$$\log_2 \mathbb{E}[d_{j,k}^2] \sim j(1 - \beta) + C \quad (3.41)$$

where C is a constant with respect to parameter j . Consider the variance estimator

$$\mu_j = \frac{1}{n_j} \sum_{k=1}^{n_j} |d_{j,k}|^2, \quad (3.42)$$

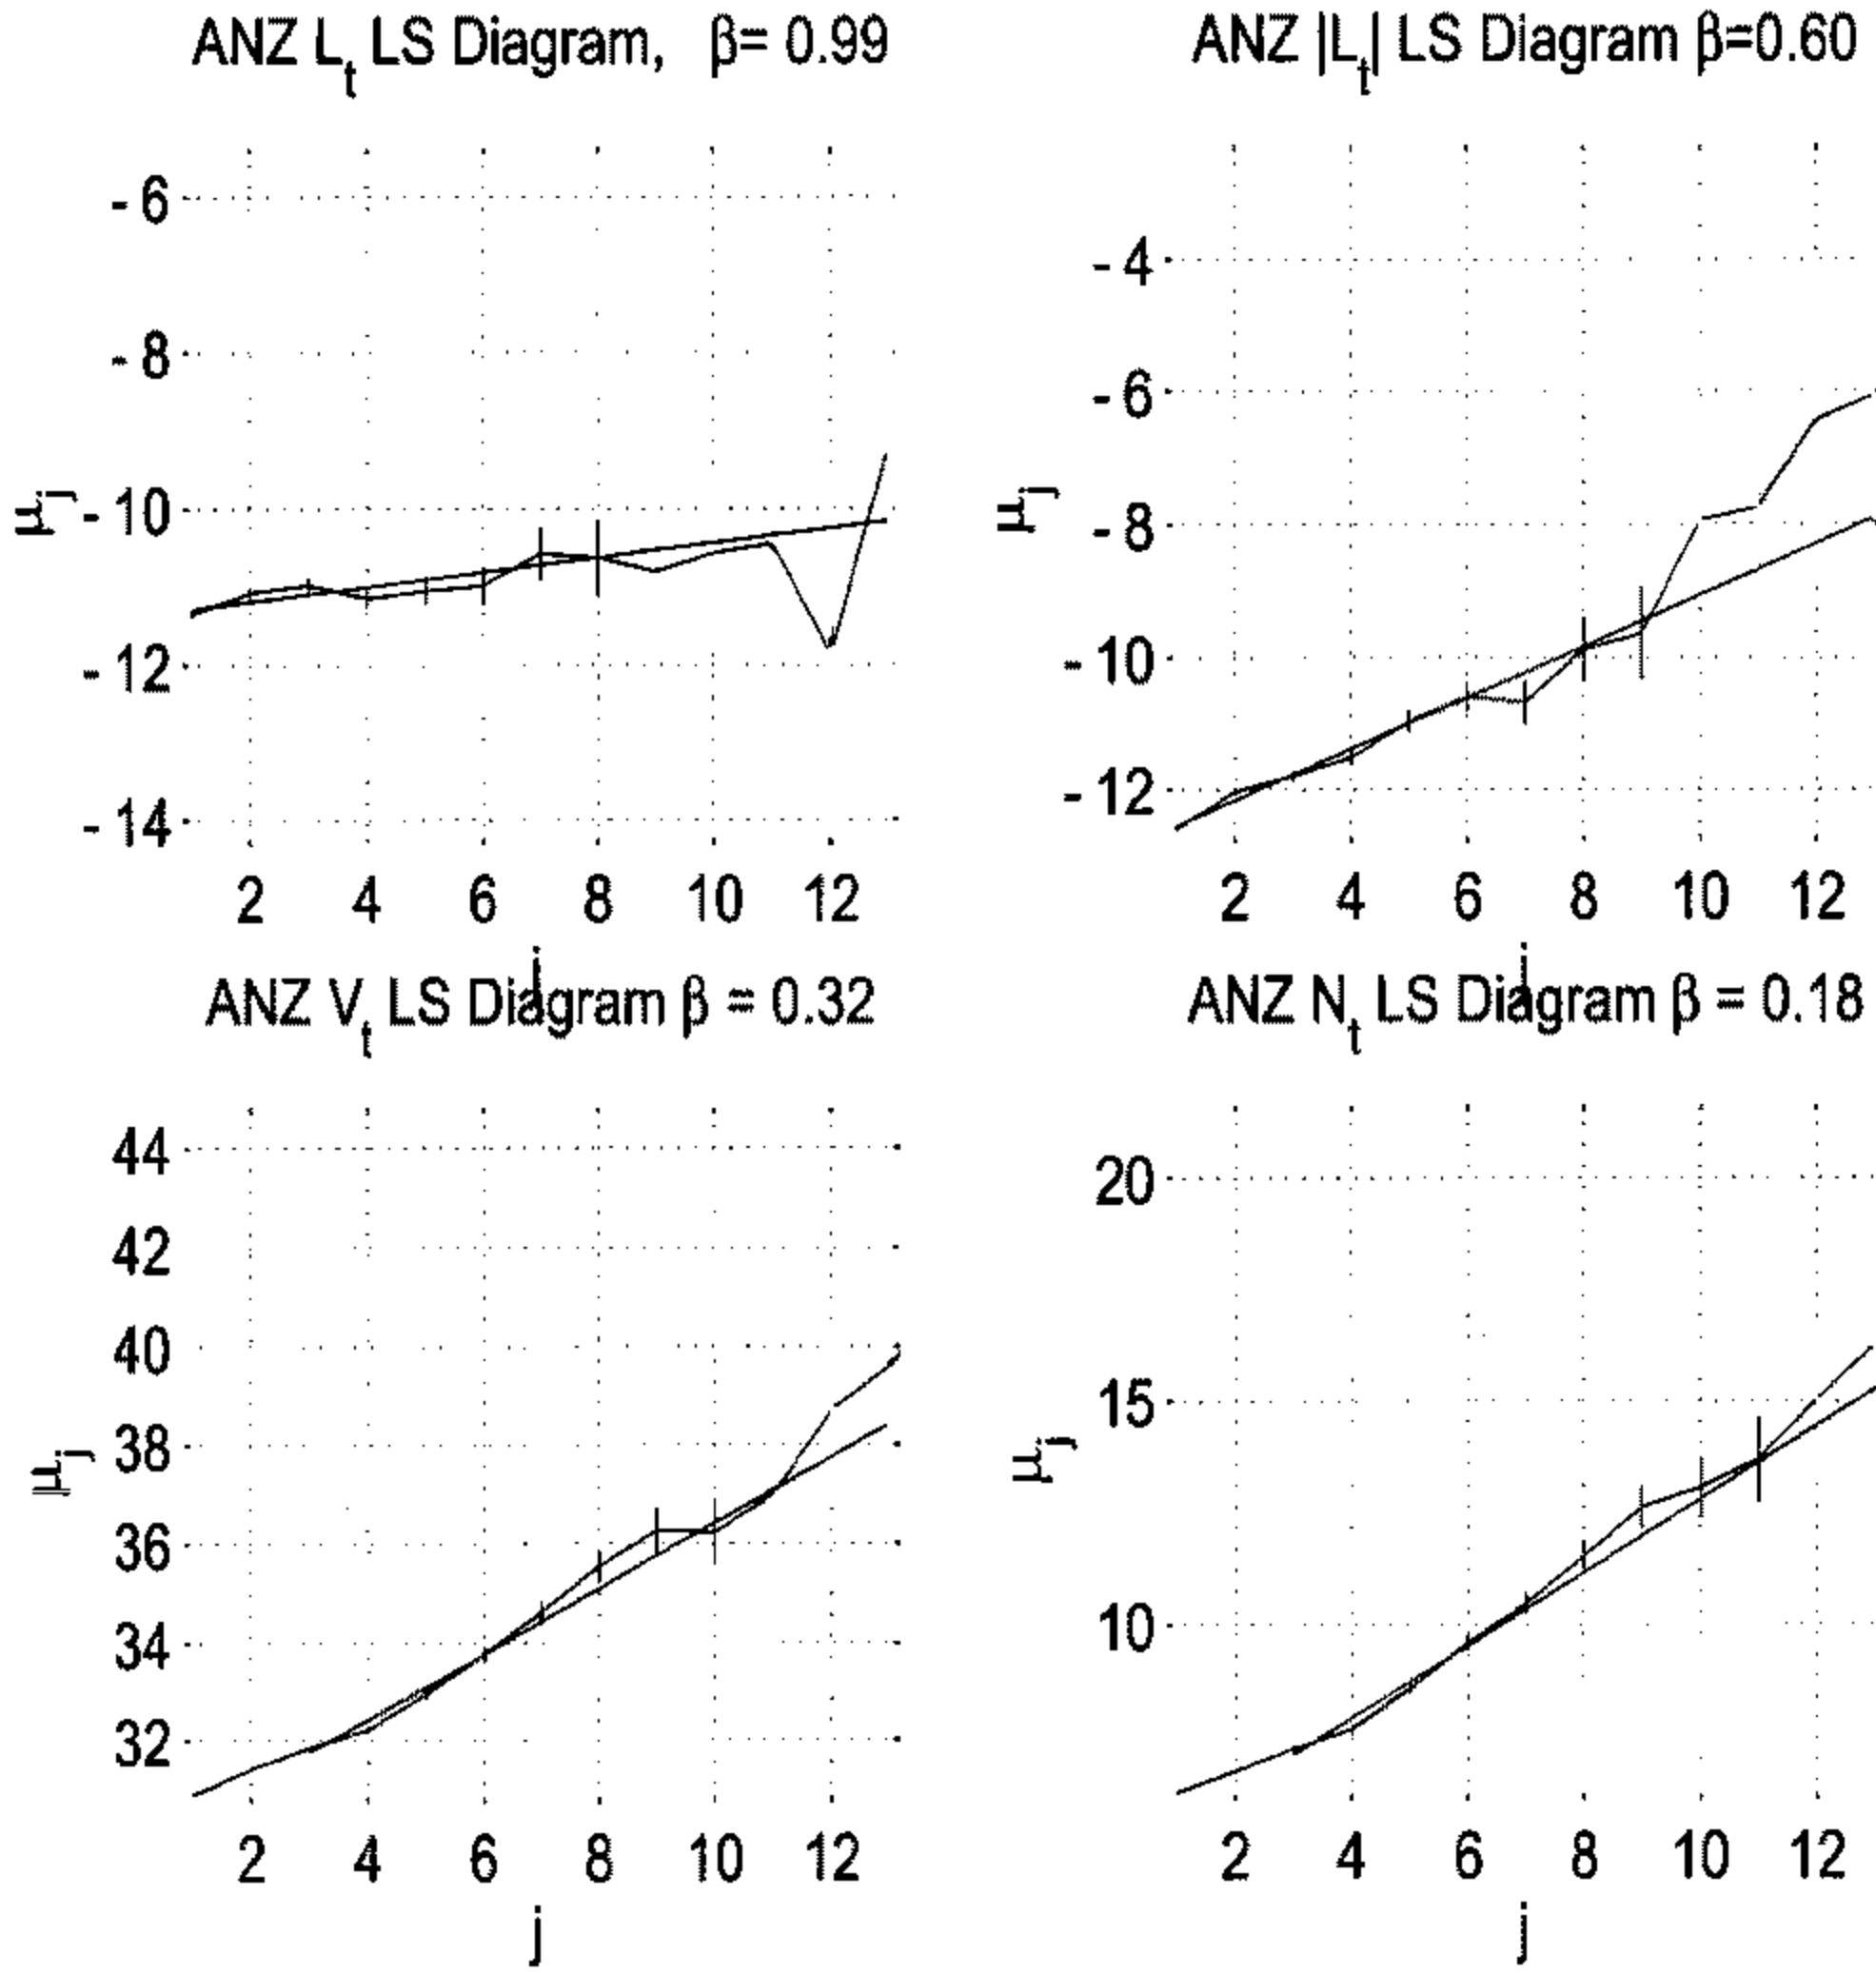


Figure 3.36: Logscale diagrams for ANZ data. The wavelet coefficients display LRD in $|L_{t,\tau}|$, $N_{t,\tau}$ and $V_{t,\tau}$.

where n_j is the number of wavelet coefficients at level j . The graph of $\log_2 \mu_i$ versus j is referred to by Abry et.al.(2000) as the logscale diagram. Estimation of the index β can be carried out via a weighted least squares fit. We used a software toolbox for MATLAB, developed by the authors of Abry et.al.(2000), (available at http://www.cubinlab.ee.mu.oz.au/~darryl/secondorder_code.html) to construct logscale diagrams and perform the weighted fits to estimate β . Included in the toolbox is a filtering routine that adapts discrete data for use with this algorithm (see Abry et.al. (2000) for details.)

In Fig. 3.36 we show the Logscale Diagram for ANZ $L_{t,\tau}$, $|L_{t,\tau}|$, $V_{t,\tau}$, and $N_{t,\tau}$. Examining the estimated values of β including those in Table B.17, we find them to be in good agreement to the previously estimated values obtained via other methods.

3.4.5 Detrended fluctuation analysis

The method of detrended fluctuation analysis (DFA) was introduced by Peng et.al. (1994) in the study of DNA nucleotides and was subsequently implemented in a financial context by Gopikrishnan et.al. (2000) and Costa and Vasconcelos (2003). This method has an advantage over the previously described variance analysis methods in that it can detect LRD in non-stationary time series. The idea is to subtract deterministic trends from the original time series and analyse the fluctuation of the detrended data.

With a time series, X_1, \dots, X_N start by forming the cumulant series,

$$y(t) = \sum_{i=1}^t X_i. \quad (3.43)$$

Next, divide $y(t)$ into $h = \frac{N}{n}$ non-overlapping intervals with n samples in each interval. The local trend in each interval, I_h , is the linear least squares fit of the data in that interval,

$$l_n(t) = a_h + b_h t \quad (3.44)$$

where a_h and b_h are the least squares coefficients in interval I_h . Define the detrended walk as,

$$y_n(t) = y(t) - l_n(t) \quad (3.45)$$

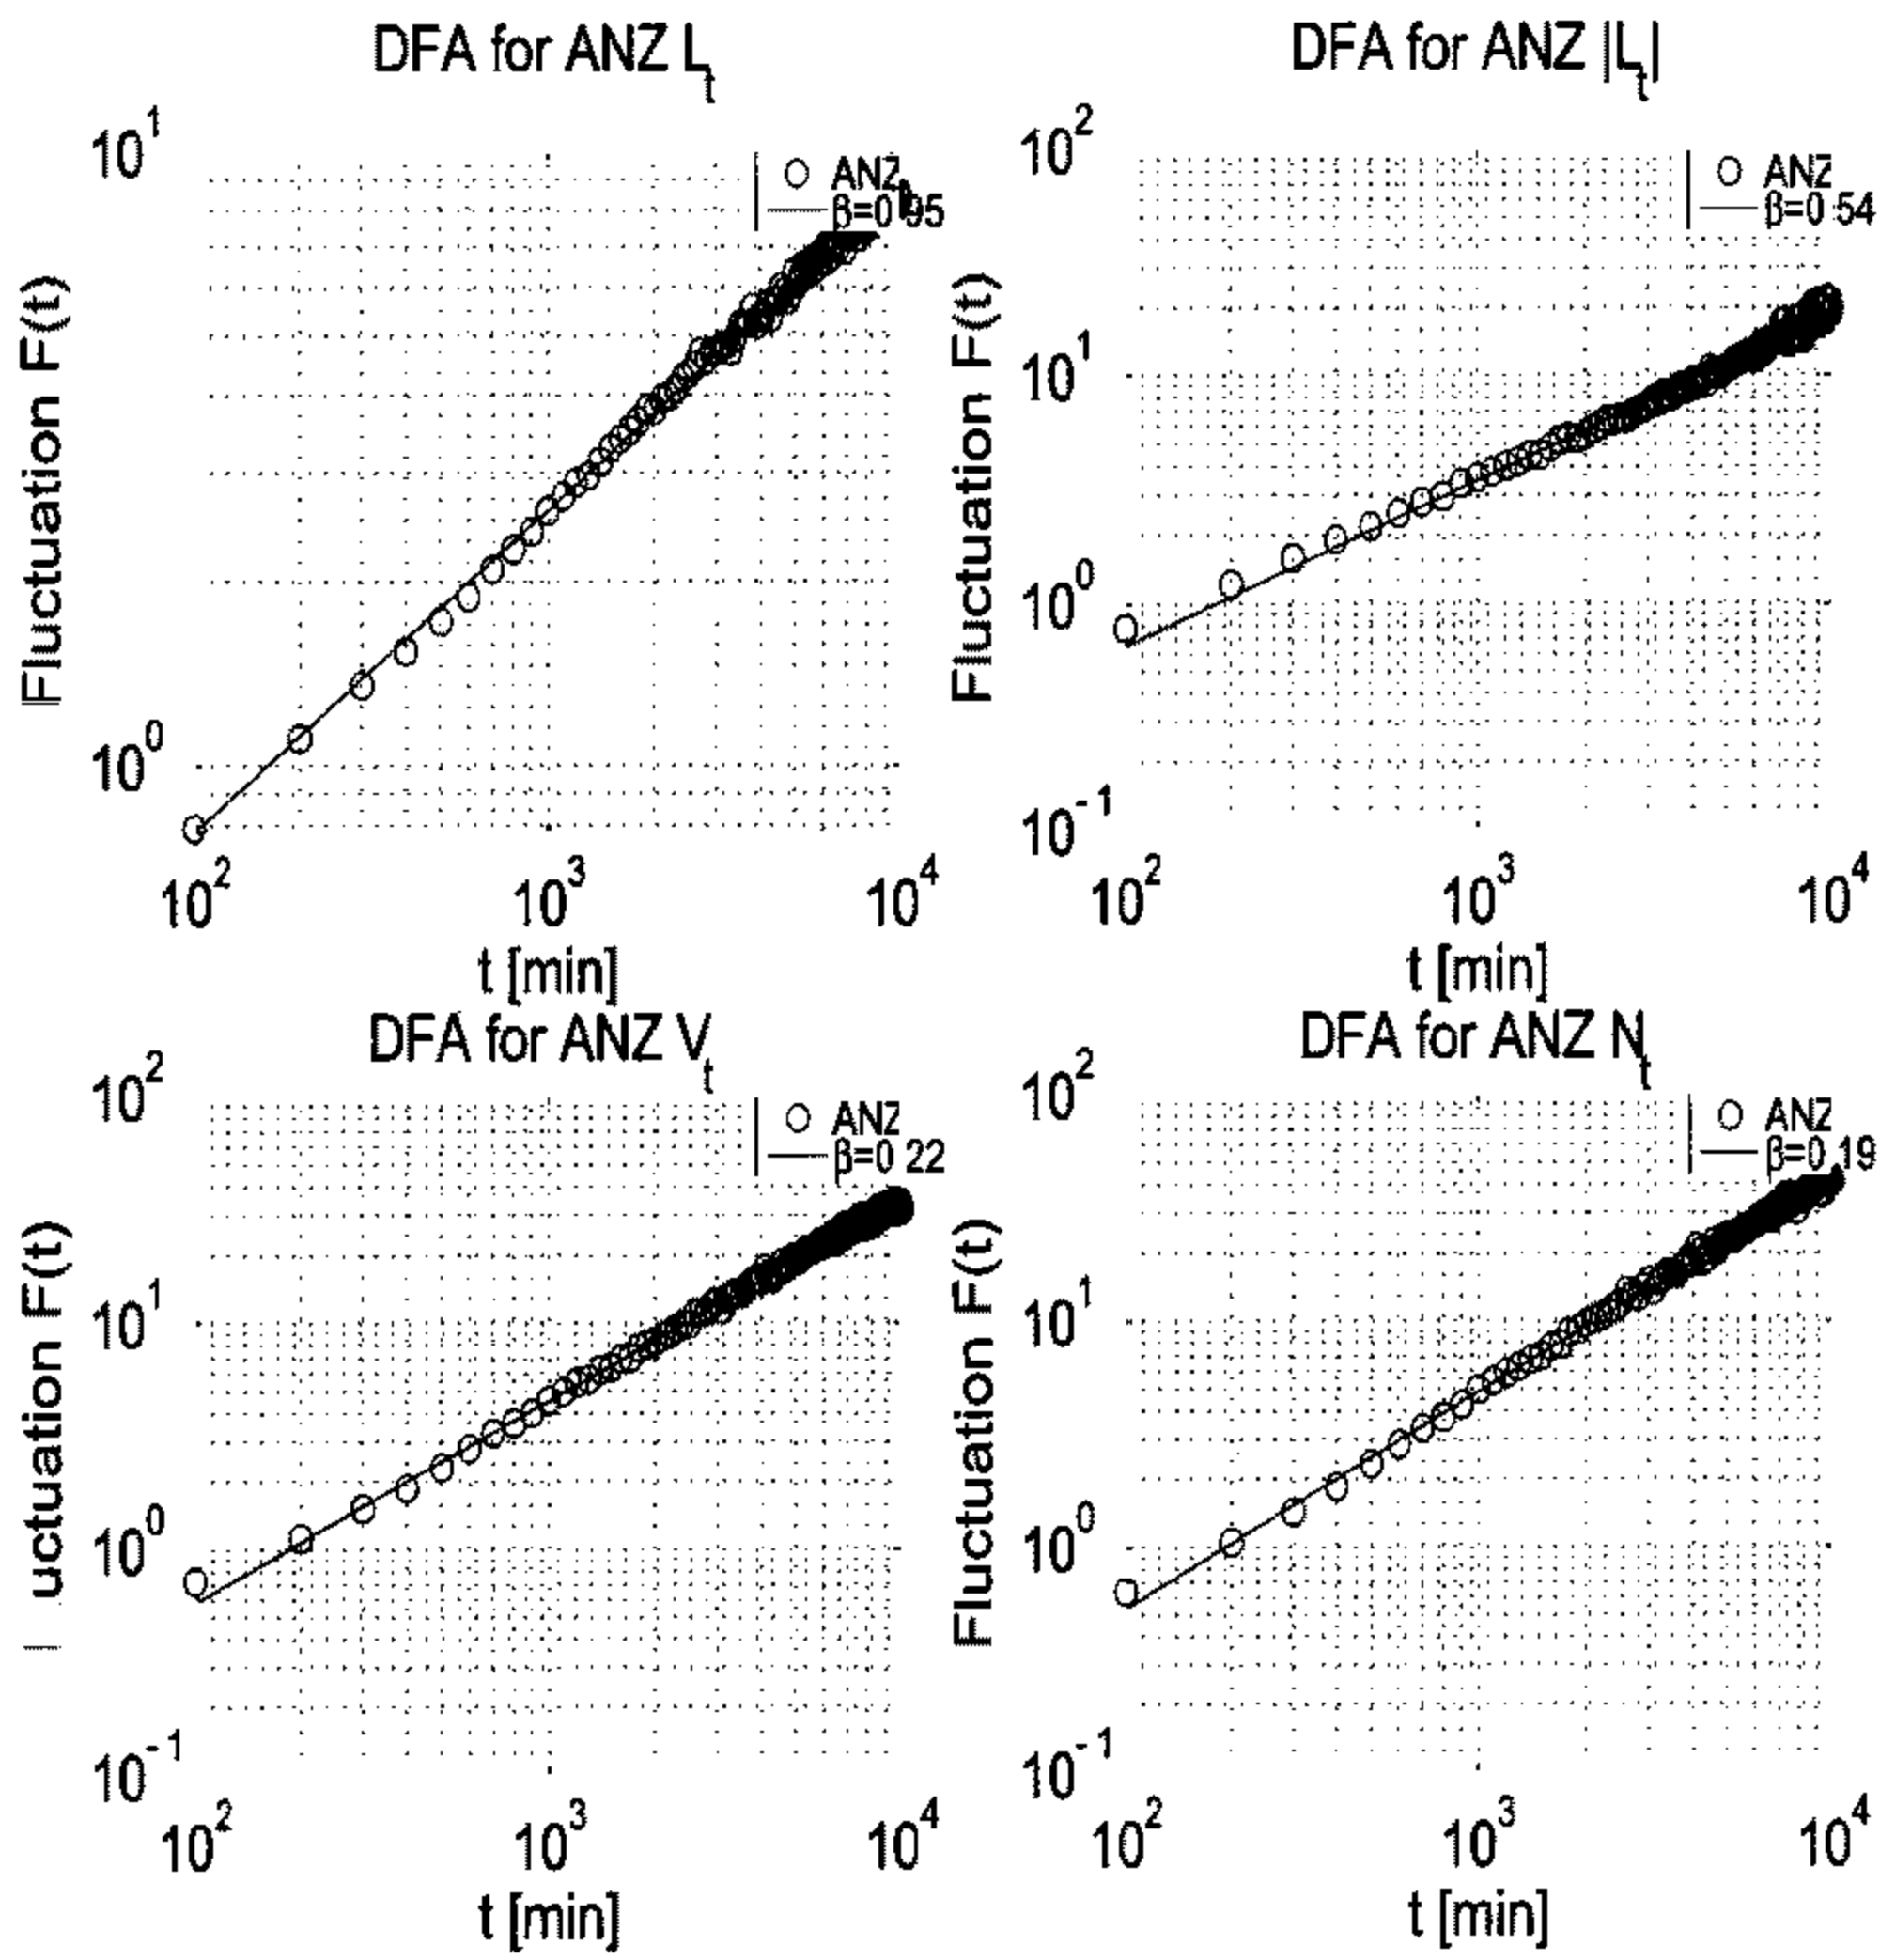


Figure 3.37: Detrended fluctuation analysis for ANZ data.

then calculate the average variance of the detrended walk for different interval lengths, n ,

$$F_d^2(n) = \mathbb{V}[y_n(t)]. \quad (3.46)$$

For a process displaying short memory the detrended walk will have the properties of a random walk so that the fluctuation function, $F_d(n)$ behaves as,

$$F_d(l) \sim n^{1/2}. \quad (3.47)$$

However, if the process possesses a power law correlation then the fluctuation function behaves as,

$$F_d(n) \sim n^{1-\beta/2}. \quad (3.48)$$

The DFA plots for ANZ data are shown in Fig. 3.37. The DFA results indicate that the log-return is uncorrelated while the absolute log-return, volume and transaction frequency all display a dependence structure that agrees with previous results. Table B.18 provides the estimates of β for other studied equities. The results agree well with the results in Tables B.13 to B.17.

3.5 Fractals and scaling

The issues of scaling and self-similarity in financial systems have been of much interest to researchers over the years. The concept of scaling in finance was born out of the ideas of geometric Brownian motion and the stable Paretian hypothesis of Mandelbrot. Under the framework of these models, price changes exhibit scale invariance. Real world data provides indications that financial markets may exhibit scaling behaviour, as has been observed when studying the power law tail behaviour of the empirical distribution function in section 3.3.3.

Scale invariance has captured the attention of many physicists who studied systems that exhibit such behaviour. Mantegna and Stanley (1995) have performed a systematic investigation of scaling behaviour of the S&P500 index. Other studies have looked at various scaling properties of other markets.

Self-similar processes have seen an increasing amount of literature devoted to their study in the last decade. A self similar process is one that obeys the property,

$$X(ct) \stackrel{d}{=} c^H X(t) \quad (3.49)$$

where H is known as the scaling or Hurst parameter. Self-similar processes are rescaled by simultaneously diluting the time axis by a factor $c > 0$ and the amplitude axis by a factor c^H . A consequence of this definition is that the moments of such a process behave as power laws of time,

$$\mathbb{E}[|X_t|^q] = |t|^{qH} \mathbb{E}[|X_1|^q]. \quad (3.50)$$

A direct result of Eq. (3.49) is that a self-similar process cannot be stationary, since the distribution of a stationary process is invariant under a time shift, *i.e.* $X(ct) \stackrel{d}{=} X(t)$ and therefore Eq. (3.49) cannot hold for any non-degenerate stationary process. Some of the self-similar processes encountered so far are:

$$\text{Wiener Process} \quad W_{ct} \stackrel{d}{=} c^{1/2} W_t \quad (3.51)$$

$$\text{fractional Brownian motion} \quad B_{ct} \stackrel{d}{=} c^H B_t \quad (3.52)$$

$$\alpha\text{-stable Lévy process} \quad L_{ct} \stackrel{d}{=} c^{\frac{1}{\alpha}} L_{t,\tau} \quad (3.53)$$

$$\text{multifractal process} \quad M_{ct} \stackrel{d}{=} c^{\tau(q)} W_t. \quad (3.54)$$

In section 3.4 we have shown that the various ASX time series have long range dependence. The notion of a self-similar process is intimately connected with that of LRD. For any self-similar process X_t with stationary increments and finite variance, the increment process $X_{t+\tau} - X_t$, will exhibit long range dependence (Eq. (3.22)) with index $\beta = 2-2H$. Also the self-similarity of $X(t)$ is transferred to its increments in the sense that the variance of the increments scales with τ ,

$$\mathbb{E}[|X_{t+\tau} - X_t|^2] = \sigma^2 |\tau|^{2H}. \quad (3.55)$$

In the following sections we will examine the scaling and fractal properties of the ASX data, estimating the scaling parameter by various methods in an attempt to shed light on this still contentious topic.

3.5.1 Scaling the empirical distribution

It was shown previously (Fig. 3.25) that the tails of the empirical distribution function displayed power law behaviour that was scale invariant. To explore this property further we calculate empirical density functions for the time intervals $\tau = 10, 20, 30, 40$ and 60 minutes. In order to observe any scaling behaviour, we normalise each time series to have unit standard deviation,

$$\hat{L}_t = \frac{L_{t,\tau}}{\sigma_{L,\tau}} \quad (3.56)$$

$$|\hat{L}_t| = \frac{|L_{t,\tau}|}{\sigma_{|L|,\tau}} \quad (3.57)$$

$$\hat{V}_t = \frac{V_{t,\tau}}{\sigma_{V,\tau}} \quad (3.58)$$

$$\hat{N}_t = \frac{N_{t,\tau}}{\sigma_{N,\tau}}. \quad (3.59)$$

By superimposing the re-scaled density estimates for the different time intervals, as shown in Fig. 3.38, Fig. 3.39, Fig. 3.40 and Fig. 3.41, we can see that the different time scales collapse to the same density function. For the ASX data, the log-return, absolute log-return, volume, and transaction frequency all exhibit scaling over the studied range of times. This scaling behaviour of the normalised data suggests that the standard deviation of each time series should scale as a function of τ . We can investigate the nature of this behaviour by considering the relationship,

$$\sigma_\tau \sim \tau^H. \quad (3.60)$$

If this equation is plotted on a log-log scale, a linear scaling relationship is indicated by a linear trend, while multiscaling would be identified by a non-linear trend. Fig. 3.42 displays the plots of the standard deviation against τ for the four series in question. In each case a definite linear trend is apparent indicating a linear scaling behaviour, at least over the time intervals $\tau = 10 - 60$ min. It is precisely this property that allows for the construction of the business time scale we used to de-seasonalise the data (section 3.2.3). Table B.19 contains the estimated slopes of the linear fit for the stocks under investigation.




Figure 3.38: Rescaled density functions for the CBA log-return, $L_{t,\tau}$. The functional form of the density appears to be preserved under scaling.




Figure 3.39: Rescaled density functions for the CBA absolute log-return, $|L_{t,\tau}|$. The functional form of the density appears to be preserved under scaling.




Figure 3.40: Rescaled density functions for the CBA absolute log-return, $V_{t,\tau}$. The functional form of the density appears to be preserved under scaling.




Figure 3.41: Rescaled density functions for the CBA absolute log-return, $N_{L,T}$. The functional form of the density appears to be preserved under scaling.

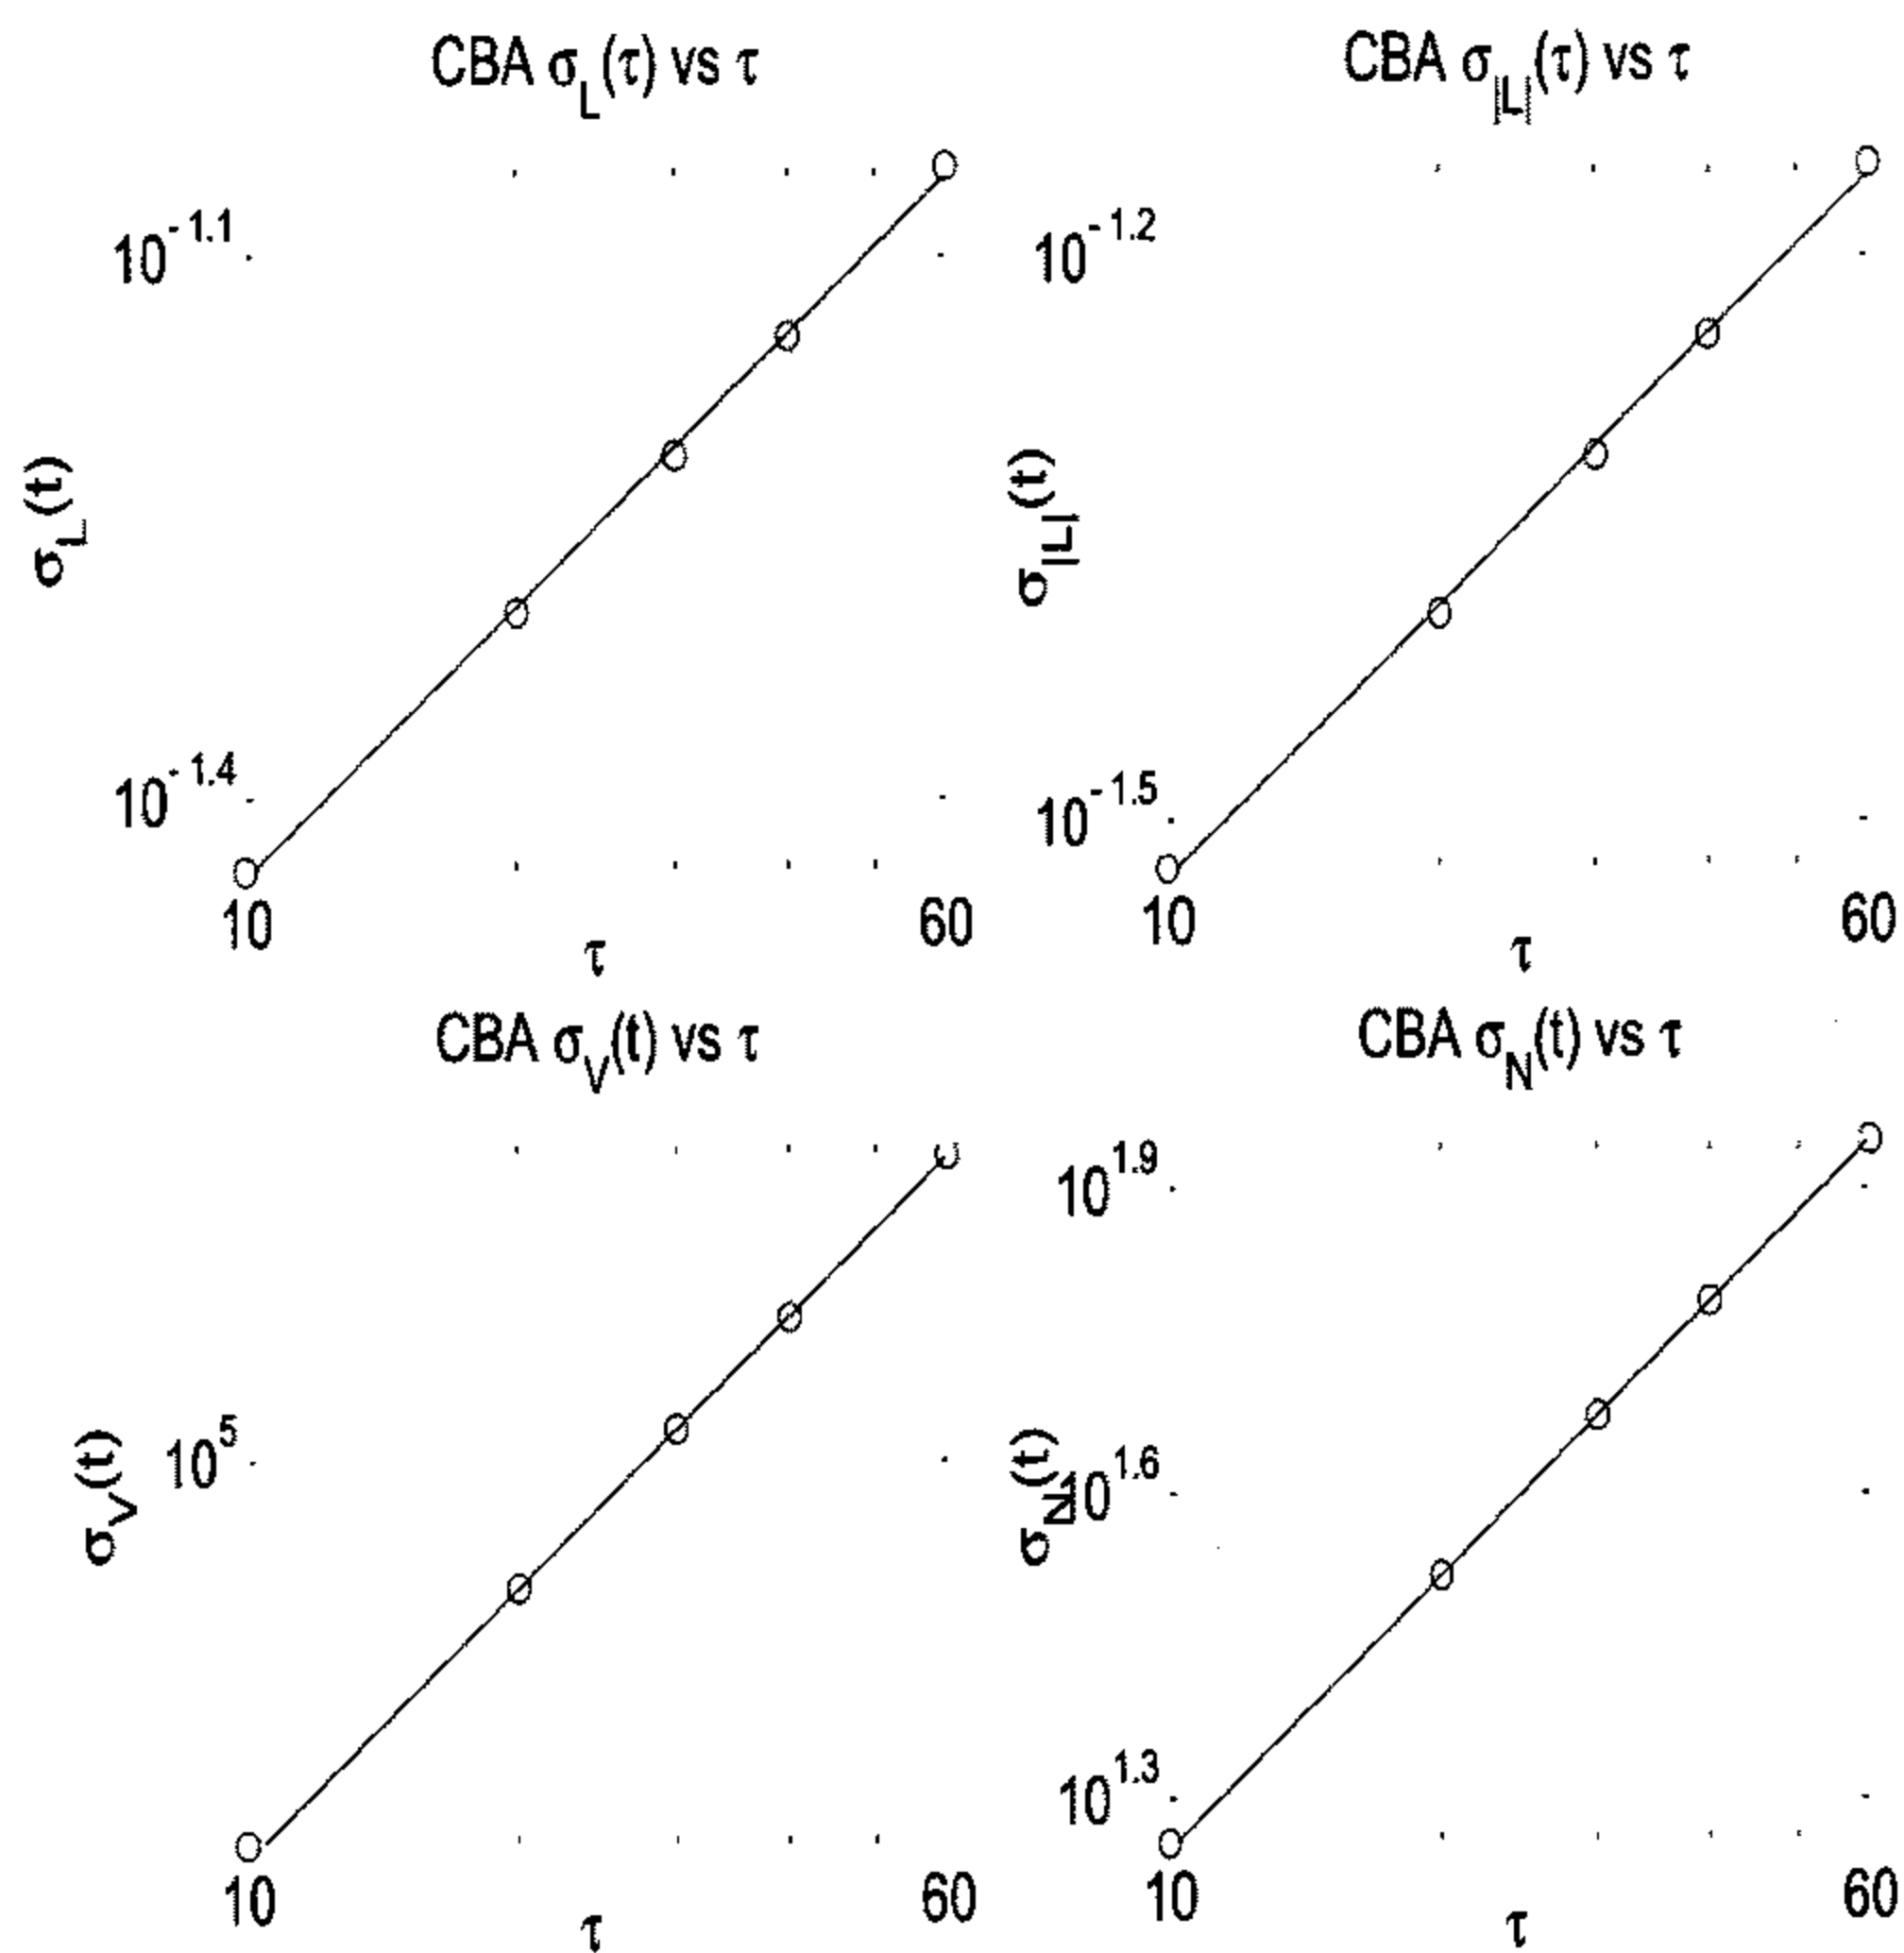


Figure 3.42: Scaling relationships for the standard deviation of each time series.

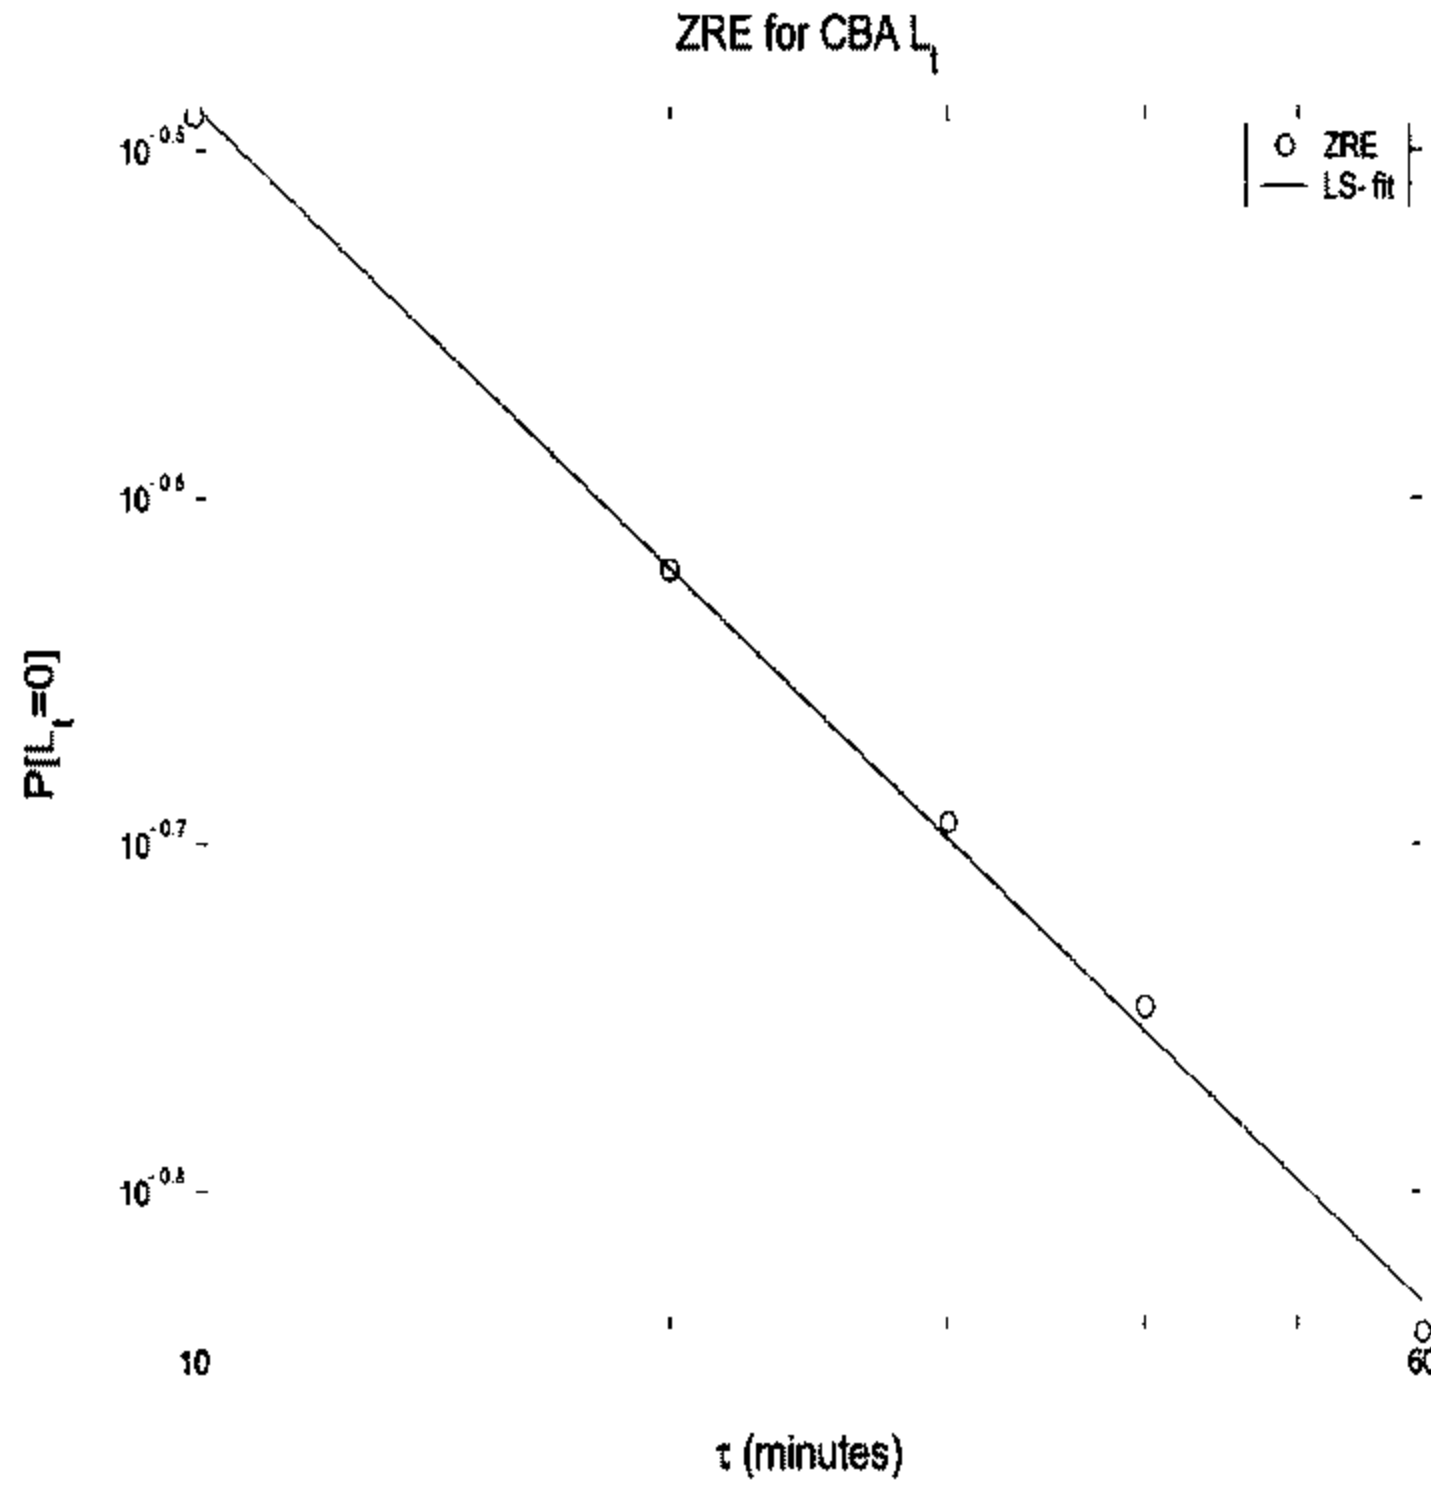


Figure 3.43: Scaling behaviour of Zero Return Enhancement, $P[L_{t,\tau} = 0]$.

Given a large enough interval size τ , this scaling must eventually breakdown as the process has finite variance with $\alpha > 2$. As shown by Bouchaud et.al. (2000) and Goprishkan et.al. (1999) there exist processes with correlation structures similar to financial data that will converge slowly to Gaussian. Evidence from other studies using low frequency data (many decades in length), such as by Goprishkan et.al.(1999), indicates that this convergence takes place on a scale of approximately one week, however our data set is too short to conduct such a study. Instead we restrict ourselves to examining high frequency scaling behaviour.

Turning our attention to the centre of the distribution we can examine the scaling of zero return enhancement. In Fig. 3.43 we have plotted the ZRE probability for time intervals, $\tau = 10, 20, 30, 40$ and 60 minutes. Viewing on a log-log scale we find that the level of ZRE scales as,

$$P[L_t = 0] \sim \tau^{-\zeta} \tag{3.61}$$

with $\zeta = 0.44$ for CBA data. Table B.20 lists the estimated values for ζ from several different stocks in the data set. We must note here that the ZRE scaling behaviour is different to the scaling behaviour of the distribution tails.

3.5.2 Scaling through the wavelet transform

As in the phenomenon of long range dependence, we can use wavelet transforms to test for self-similarity. It has been shown by Abry et.al. (1997) that the wavelet coefficients of a self-similar process replicate the self-similarity property. Suppose that X_t is a self-similar process with parameter H then for fixed $j \in \mathbb{Z}$,

$$d_X(j, k) \stackrel{d}{=} 2^{j(H+\frac{1}{2})} d_X(0, k). \quad (3.62)$$

Thus the moments will follow,

$$\mathbb{E}[|d_X(j, k)|^q] = 2^{jq(H+\frac{1}{2})} \mathbb{E}[|d_X(0, k)|^q]. \quad (3.63)$$

In addition to this, when X_t possesses stationary increments,

$$\mathbb{E}[|d_X(j, k)|^2] = C 2^{j(2H+1)} \quad (3.64)$$

with $C = \mathbb{E}[|d_X(0, 0)|^2]$. Taking logarithms of both sides yields the linear relationship,

$$\log_2 \mathbb{E}[|d_X(j, k)|^2] = (2H + 1)j + C' \quad (3.65)$$

where the scaling parameter is obtained from the gradient $(2H + 1)$. This estimator takes the same form as that of Eq. (3.41) and serves to illustrate the close connection between self-similarity and LRD, with both sharing a fundamental power law behaviour in the wavelet domain.

As for LRD, we test for self-similarity with the logscale diagram, using the MATLAB toolbox of Abry et.al. (2000) to measure,

$$\mu_i = \frac{1}{n_j} \sum_{k=1}^{n_j} |d_{j,k}|^2. \quad (3.66)$$

Recalling that a self-similar process cannot be stationary, we will focus our attention on the log-price,

$$Q_L(t) = \sum_{i=1}^t L_i \quad (3.67)$$

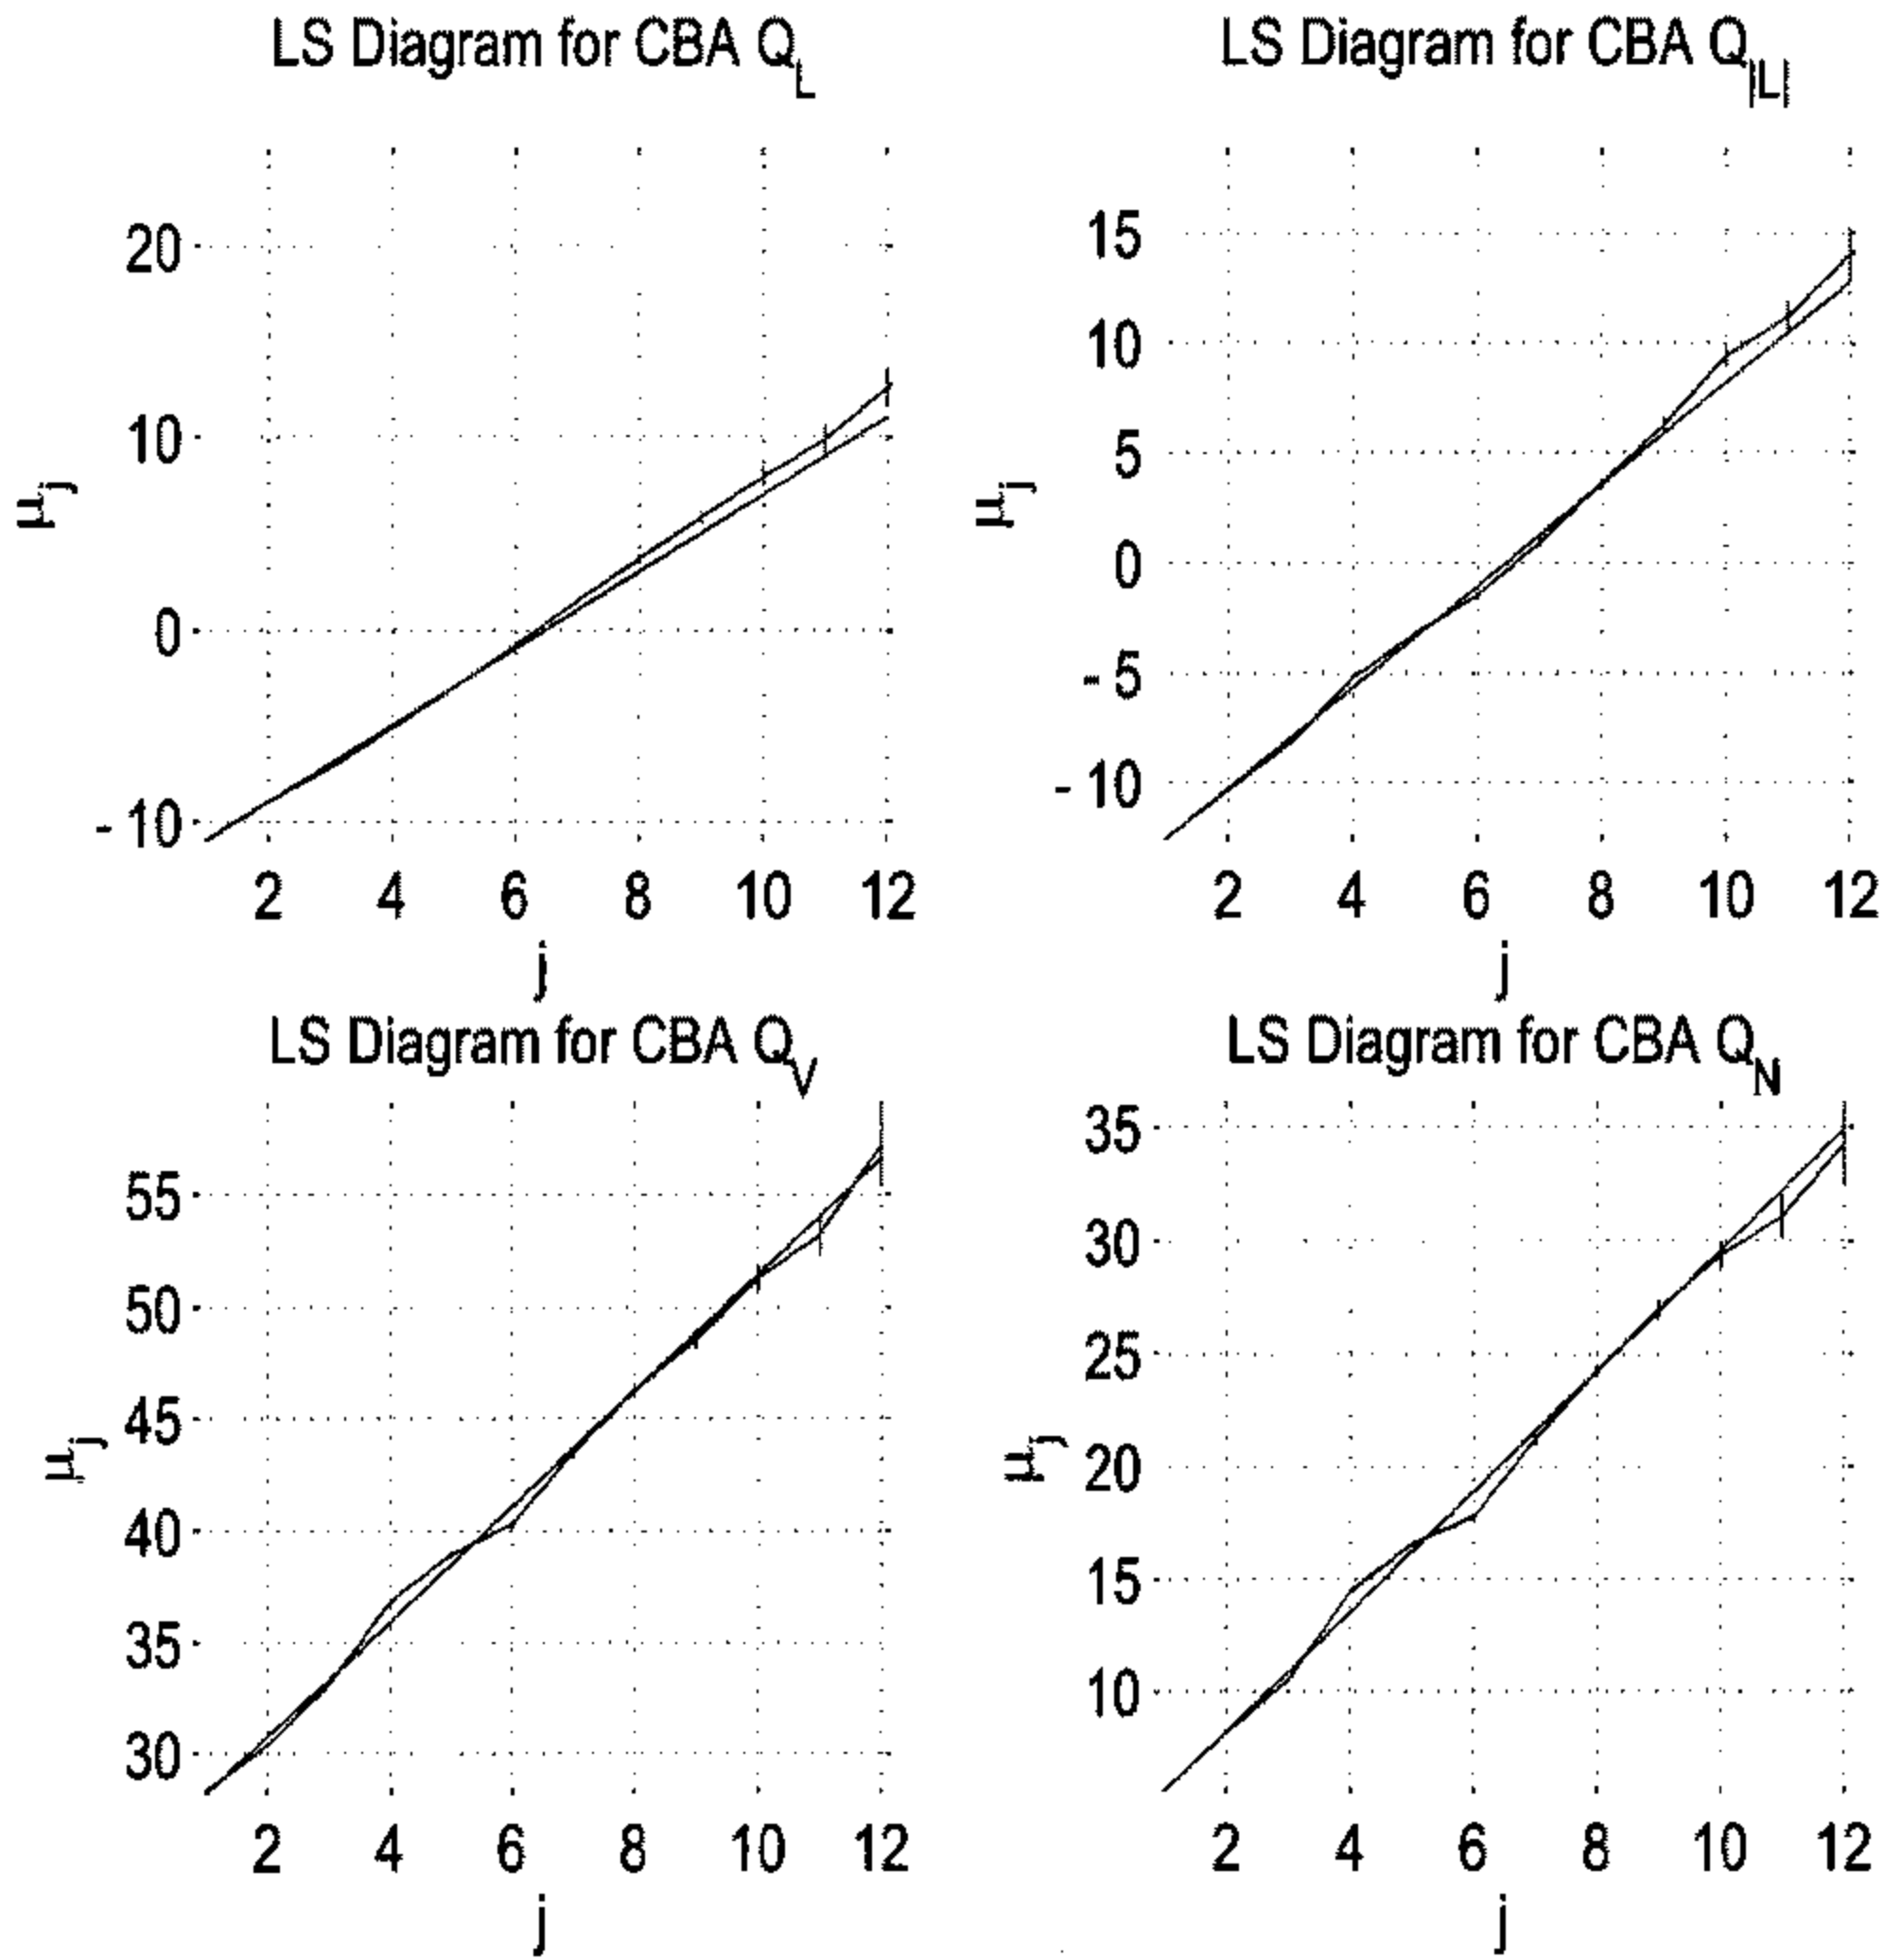


Figure 3.44: Logscale digram for the self-similarity of cumulant series, $Q_L(t), Q_{|L|}(t), Q_V(t)$ and $Q_N(t)$.

and the cumulant series of the absolute return, volume, transaction frequency,

$$Q_{|L|}(t) = \sum_{i=1}^t |L_i| \quad (3.68)$$

$$Q_V(t) = \sum_{i=1}^t V_i \quad (3.69)$$

$$Q_N(t) = \sum_{i=1}^t N_i. \quad (3.70)$$

These cumulative processes represent the random walks undertaken by the various financial processes. Fig. 3.44 are the logscale diagrams for the CBA cumulant series (Eqs. (3.67-3.70)). The estimated values for the scaling parameter H are in agreement to the empirically measured scaling values of Section (3.5.1). Table B.21 summarises the logscale diagram estimates for the data. Using $\beta = 2 - 2H$ which relates LRD to self-similarity we find that measured values of H and β (Tables (B.13-B.18)) are in good agreement with each other.

3.5.3 Examining the fractal nature

There has been a growing interest in fractal behaviour of financial markets in recent years. Researchers such as Heyde (1999) and Mandelbrot (1997) have suggested fractal and multifractal behaviour of financial data. In addition Gorski et.al. (2002) and Ho et.al. (2004) have undertaken studies into this type of behaviour. To investigate the fractal nature of the ASX equity data we employ the box counting exponent of Renyi (1970). This method provides a way to visually identify not just the presence, but also the type of scaling behaviour in the data.

The box counting method is the basic paradigm for the practical calculation of the generalised Renyi entropy for a function $X(t)$ of length N . This method consists of superimposing a grid of N_δ boxes, each of length δ , onto the spatial pattern given by X_t . For each box $1 < i < N_\delta$, count the number of points, n_i , in each box. Write the mass of the box as,

$$p_i = \frac{n_i}{N}. \quad (3.71)$$




Figure 3.45: The above plots show the estimated Renyi exponents and generalised dimension D_q for ANZ $Q_L(t)$. $q = 0.5, 0.9, 1.5, 2, 2.5, 3$ for circles, plus signs, triangles, stars, diamonds and squares respectively. Solid lines represent least squares fits.




Figure 3.46: The above plots show the estimated Renyi exponents and generalised dimension D_q for ANZ $Q_{|L|}(t)$. $q = 0.5, 0.9, 1.5, 2, 2.5, 3$ for circles, plus signs, triangles, stars, diamonds and squares respectively. Solid lines represent least squares fits.




Figure 3.47: The above plots show the estimated Renyi exponents and generalised dimension D_q for ANZ $Q_V(t)$. $q = 0.5, 0.9, 1.5, 2, 2.5, 3$ for circles, plus signs, triangles, stars, diamonds and squares respectively. Solid lines represent least squares fits.




Figure 3.48: The above plots show the estimated Renyi exponents and generalised dimension D_q for ANZ $Q_N(t)$. $q = 0.5, 0.9, 1.5, 2, 2.5, 3$ for circles, plus signs, triangles, stars, diamonds and squares respectively. Solid lines represent least squares fits.

The generalised entropy (Renyi 1970) is defined as,

$$I_q(\delta) = \frac{1}{1-q} \log \sum_{i=1}^{N_\delta} p_i^q \quad (3.72)$$

such that the generalised dimension, D_q for the q^{th} fractal moment is given by,

$$D_q = - \lim_{\delta \rightarrow 0} \frac{I_q(\delta)}{\log \delta}. \quad (3.73)$$

In practice the value of D_q can be estimated by plotting $I_q(\delta)$ against $\log \delta$. For a simple fractal object, D_q is a linear function of q , as no extra information will be gained by examining higher order moments. However, for a multifractal object, D_q will display non-linear behaviour in q , that requires an infinite number of scaling exponents in its description.

The Renyi box counting algorithm is performed on $Q_L(t)$, $Q_{|L|}(t)$, $Q_V(t)$, and $Q_N(t)$. Fig. 3.45, Fig. 3.46, Fig. 3.47 and Fig. 3.48 shows plots of $I_q(\delta)$ vs. $\log \delta$ and D_q vs. q . The results of these calculations indicate a linear behaviour for D_q as a function of q . Thus over the time scales $\tau = 10$ to 60 minutes ASX financial data appears to display fractal properties. We might expect to see a non-linear trend corresponding to the slow break down of scaling over larger time scales however this behaviour is too slight to be observed in these estimates. Hence, to a good approximation, one may consider a fractal scaling relationship sufficient to describe high frequency data.

3.6 Summary

In this chapter we have carried out an empirical investigation of high frequency equity data for the Australian Stock Exchange over the period January 1993 to July 2002. We examined the return series, its absolute value, the volume traded and the transaction frequency for both previously reported and unreported behaviour. This behaviour was compared with that found in other studies. The properties of the data that were uncovered in this chapter provide insight into the stochastic processes that drive financial markets.

It was demonstrated how, due to the time-zone of the Australian market, the ASX returns can be represented as two separate processes for the overnight and intraday periods. Further, we have shown that while the intraday log-return itself contains no seasonal trend, the absolute log-return, volume and transaction frequency all display strong periodic behaviour. This periodic behaviour is due in part to microstructure effects that are unique to the ASX. We would not expect studies on different exchanges to yield similar results, as no two exchanges operate under the same conditions. The extent to which market regulations affect the results of the commonly performed analysis can only be determined by a wider study across several markets.

The performance of previously proposed models for stock prices was investigated using high frequency log-return data. The models of Mandelbrot and Taylor (1967), Praetz (1972), Madan and Seneta (1990) and Mantegna and Stanley (1998) fit the data and were compared with geometric Brownian motion. The log-likelihood values for the fitted distributions we found that a model of Madan and Seneta (1990) best describes the high frequency data, followed by the Student- t model of Praetz (1972), which performed best using the truncated Lévy distribution and the lognormal distribution and GBM were the worst performers. The conditional dependence structure of the data is not fully represented by a GARCH(1,1) model's autocorrelation function decaying too rapidly. The tail behaviour of the empirical distributions of the data and returns appear to possess power law type behaviour, consistent with other studies. The estimated power law exponents were found to differ from those presented in the literature on other markets. In particular, the estimated volume was found to possess a tail index of more than twice that of the returns (see Gabaix et.al. (2003a), Gabaix et.al. (2003b) and Gopikrishnan et.al. (2003)). The relative values of the estimated tail indices for L_t , $|L_{t,\tau}|$, $V_{t,\tau}$ were found to differ from those found in previous studies, making the ASX returns unique with the explanation of the so-called cubic and half-cubic laws as

proposed in Gabaix et.al. (2003a). It remains to be seen if the different power law behaviour found on the ASX is market and or regional specific.

The correlation structure of the ASX data was investigated for long memory behaviour using several methods. By examining the autocorrelation function, the spectral density and using the methods of variance plots, R/S analysis, wavelet transforms and detrended fluctuation analysis, we were able to obtain consistent estimates of long range dependence in the data. Results showed that while the log-return process appeared to be uncorrelated with an exponent $\beta_{L_t} \approx 1$ the absolute log-return, volume and transaction frequency all displayed behaviour consistent with that of a long memory process.

Finally, we explored the scaling behaviour of the ASX data. Scaling properties of financial data are of much interest to researchers and are linked to the phenomena of LRD. By rescaling the ASX data for different sampling intervals we showed that the probability density functions collapsed to the same functional form. Plotting the variance as a function of interval size we obtained rough estimates of the scaling index, H of the data. We examined the cumulants of the various time series in detail, using methods based on the wavelet transform to estimate H . These results are found to be in agreement and to correspond with the LRD estimates of Section 4. The data was explored further by calculating the Renyi exponents in order to test for possible multifractal behaviour. We found that, to a good approximation, ASX high frequency data displays only fractal properties. This fractal behaviour explains why the power law exponents of Section 3.3 do not seem to change with τ , shown by Fig. 3.25.

Chapter 4

New approaches for price modelling

In the previous chapter we presented the main defining properties of the stochastic nature of Australian financial markets. It was also shown that past models had various degrees of difficulty in describing certain aspects of high frequency data. In particular the previous models focused on proposing specific distributions to describe the log-return. These models all lacked a distinct reasoning as to why these distributions should model the log-return. Subordinated models partially address this problem by assuming that the price process develops as a GBM over a random time scale. However, this approach still reduces to speculating distributions almost arbitrarily. More recent trends in financial research have focused upon describing specific features of financial markets.

The recent rise in the availability of high frequency financial data has seen an increase in the number of studies focusing on the areas of classification and modelling of financial markets. The development of models that are able to reflect the various observed phenomena of real data is an important step towards obtaining a full understanding of the fundamental stochastic processes that drive the market. Previous studies such as Gabaix et.al. (2003a), Mandelbrot (1963), Costa and Vasconcelos (2003) and Stanislavsky (2003) have proposed models in order to explain the existence and origin of various market properties including: fat-tailed distributions;

power law relationships in the log-return, volume and transaction frequency; temporal correlation and long memory of volatility; scaling properties of the log-return and non-stationary behaviour of financial time series. Models seeking to describe these phenomena have been developed from behaviour reported by recent empirical studies on high frequency financial data such as Gopikrishnan et.al. (2000), Molgedey and Ebeling (2000), Antoniou et.al. (2004), Ho et.al. (2004), Mizuno et.al. (2003) and Gorski et.al. (2002).

This chapter continues the theme of investigating the properties of financial markets through modelling, offering a possible explanation of the origin of the large number of zero returns found in equity data from the Australian Stock Exchange. We investigate the existence of zero return enhancement, a phenomenon previously reported in the DAX index Gorski et.al. (2002), and present a threshold model that is able to describe the effect. This model can explain phenomenologically the disproportionate number of zero returns and how it is related to the activity level of the market place. With our model we are able to demonstrate the effects that small scale zero return enhancement has on large scale stock returns.

The continuous time random walk has recently been proposed as a model for high frequency stock prices by Scalas et.al. (2000). This model takes a phenomenological approach rather than relying on the efficient market hypothesis and can operate in a non-Markovian framework. We revisit the theory of continuous time random walks and the model of Scalas et.al. (2000). Using the ASX high frequency data we contribute to the development of the model by testing its assumptions and behaviour. Much of this chapter appears in published form in Bertram (2005).

4.1 Threshold Brownian motion

We adopt a phenomenological approach to asset modelling in the sense that we want our models to be able to describe some specific phenomena observed in the data. It is desirable to be able to apply some reasoning when proposing specific processes or distributions to describe the data, as this will lead to a better understanding of the process of price change. In an ideal market, where market activity and volatility are

constant in time and traders are all equally informed, it seems reasonable that price should follow a geometric Brownian motion. Hence, log-returns should be uncorrelated and normally distributed. In reality, financial markets are far from ideal. As we saw in Chapter 3, there are various discrepancies and irregularities within the market. For instance, the level of market activity varies greatly throughout the day. We also saw that traders experience time intervals where there is a lack of incentive to change the price. This leads to the effect of zero return enhancement.

The model we present in this section represents a first attempt to describe the occurrence of zero returns of type 2 and type 3, as defined in section 3.2.4. According to the efficient market hypothesis the price of a stock reflects the total information available to the marketplace. Hence, a change in the level of information leads to a price change. Based on the idea of identifying price change with information levels we make the following statement concerning time intervals where the price has not changed, causing a zero return.

A zero return corresponds to a time interval where the amount of information, or market activity level, has been insufficient to affect a price change.

Owing to the low level of market activity at these times, traders have insufficient incentive to change the price of the stock and simply continue to trade at the previous price. Periods such as these may arise many times during the course of a trading day, strongly contributing to the dynamics of the stock price.

We propose the following model to describe the zero return enhancement features of the ASX taking into account the varying levels of information present in the market. Consider the following stochastic differential equation to describe the price change over time interval dt ,

$$dp_t = p_t[\mu dt + \sigma dW_t] \mathbf{1}(A_t > a) \quad (4.1)$$

where μ and σ^2 are the instantaneous mean and variance of the stochastic process driving the random price change and

$$\mathbf{1}(A_t > a) = \begin{cases} 1 & A_t > a \\ 0 & A_t \leq a. \end{cases}$$

Depending on the values of A_t and a , the stock price described by Eq. (4.1) will evolve according to a geometric Brownian motion, or it will remain constant. We may think of A_t as an activity process representing the random amount of information available to traders at time t , and a as a threshold level below which there is insufficient activity to affect a price change. We call the model threshold Brownian motion (TBM). This new model extends the classic GBM model but it can be further generalised by replacing the GBM with an alternative model for random price fluctuation. In the case of Eq. (4.1) we have used GBM in order to retain some analytic tractability.

Over a small time interval the model should display a marked departure from Gaussian but as the time scales increases, the contribution of the indicator function will die away slowly and the process will become Gaussian.

4.1.1 The model

Our model is similar in structure to the many subordinated and compound models proposed to describe the fat-tailed nature of the log-return distribution. If we assume that the activity process A_t is independent of the Wiener process driving the stock price fluctuations, W_t , the solution of Eq. (4.1) can be written as,

$$p_t = p_0 e^{(\mu - \frac{1}{2}\sigma^2) \int_0^t \mathbf{1}(A_s > a) ds + \sigma \int_0^t \mathbf{1}(A_s > a) dW_s}. \quad (4.2)$$

Using the representation in Eq. (4.2) we can formulate an expression for the log-return, $L_{t,\tau}$, over time interval τ ,

$$L_{t,\tau} = (\mu - \frac{1}{2}\sigma^2) \int_t^{t+\tau} \mathbf{1}(A_s > a) ds + \sigma \int_t^{t+\tau} \mathbf{1}(A_s > a) dW_s. \quad (4.3)$$

To explore this model further we will examine the distribution properties of the log-return. The distribution of the log-return can be derived via conditioning, as described in Section 2.3.1. The probability density function of the log-return over time interval τ , $f_{t,\tau}(x) = \mathbf{P}[L_{t,\tau} \in dx]$, is given by,

$$f_{t,\tau}(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_0^\tau \frac{1}{\sqrt{\nu}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\nu)^2}{2\sigma^2\nu}} dG_{t,\tau}(\nu) \quad (4.4)$$

where $G_{t,\tau}(\nu) = \mathbf{P}[T_{t,\tau} \leq \nu]$ is the distribution function of the random variable,

$$T_{t,\tau} = \int_t^{t+\tau} \mathbf{1}(A_s > a) ds. \quad (4.5)$$

The variable $T_{t,\tau}$ represents the length of time spent above the threshold level a during time interval $[t, t + \tau]$ and is commonly referred to as the occupation time of the process A_t . Taking the Fourier transform in x of Eq. (4.4) gives the characteristic function of the log-return process,

$$\phi_{t,\tau}(\theta) = \int_0^\tau e^{i(\mu - \frac{1}{2}\sigma^2)\nu\theta - \frac{1}{2}\sigma^2\nu\theta^2} dG_{t,\tau}(\nu). \quad (4.6)$$

Recall from Section 2.3.1 that we can derive expressions for the k^{th} moment, m_k , of the log-return distribution using the formula,

$$m_k = \mathbb{E}_L[L_{t,\tau}^k] = (-i)^k \left. \frac{\partial^k \phi_{t,\tau}}{\partial \theta^k} \right|_{\theta=0} \quad (4.7)$$

where $\mathbb{E}_L[\cdot]$ represents the expectation with respect to the distribution function of the variable $L_{t,\tau}$. Using Eq. (4.6) and Eq. (4.7), the first four moments are given by,

$$m_1 = (\mu - \frac{1}{2}\sigma^2)\mathbb{E}_T[\nu] \quad (4.8)$$

$$m_2 = \sigma^2\mathbb{E}_T[\nu] + (\mu - \frac{1}{2}\sigma^2)\mathbb{E}_T[\nu^2] \quad (4.9)$$

$$m_3 = 3\sigma^2(\mu - \frac{1}{2}\sigma^2)\mathbb{E}_T[\nu^2] + (\mu - \frac{1}{2}\sigma^2)^3\mathbb{E}_T[\nu^3] \quad (4.10)$$

$$m_4 = 3\sigma^4\mathbb{E}_T[\nu^2] + 6\sigma^2(\mu - \frac{1}{2}\sigma^2)^2\mathbb{E}_T[\nu^3] + (\mu - \frac{1}{2}\sigma^2)^4\mathbb{E}_T[\nu^4] \quad (4.11)$$

where the expected values of ν , ν^2 , ν^3 and ν^4 are the first four moments of the random variable $T_{t,\tau}$.

Turning our attention to the dependence structure of the log-return process, we can derive expressions for the covariance of the log-return, absolute log-return and squared log-return. In the case where the random stock increment process has zero drift, *i.e.* $\mu = \frac{1}{2}\sigma^2$, we have

$$\text{cov}[L_{t,\tau}, L_{t,\tau}] = 0 \quad (4.12)$$

$$\text{cov}[|L_{t,\tau}|, |L_{t,\tau}|] = \frac{2\sigma^2}{\pi} \text{cov}[\sqrt{T_{t,\tau}}, \sqrt{T_{t,\tau}}] \quad (4.13)$$

$$\text{cov}[L_{t,\tau}^2, L_{t,\tau}^2] = \sigma^4 \text{cov}[T_{t,\tau}, T_{t,\tau}]. \quad (4.14)$$

This result illustrates how the behaviour of higher order correlations of the process $L_{t,\tau}$ follow from the correlation structure of the occupation time variable $T_{t,\tau}$.

The above results display the extent to which the model's behaviour depends on the properties of the occupation time distribution, $G_{t,\tau}(\nu)$. This distribution is itself determined by the choice of activity process A_t .

4.1.2 Occupation time

The occupation time of a function $X(t)$ is defined as being the total amount of time spent above some threshold level, a , during the time interval $[0, t]$. Mathematically, this problem can be stated with the following integral representation,

$$T(t) = \int_0^t \mathbf{1}(X(s) > a) ds \quad (4.15)$$

where $\mathbf{1}(x > a)$ is the indicator (or step) function given by,

$$\mathbf{1}(x > a) = \begin{cases} 1 & x > a \\ 0 & x \leq a. \end{cases} \quad (4.16)$$

Evaluating the above integral delivers a value, $T(t)$, for the amount of time spent above level a during $[0, t]$. Fig. 4.1 illustrates the occupation time for the function $X(t) = \sin(t)$ with threshold level, $a = \frac{1}{2}$.

A similar situation exists in stochastic calculus, where the deterministic function $X(t)$, is replaced by a stochastic process X_t , giving the integral,

$$T_{0,t} = \int_0^t \mathbf{1}(X_s > a) ds. \quad (4.17)$$

In this case the occupation time, $T_{0,t}$, is a random variable. Hence, the quantity of interest is not the occupation time itself but rather its probability density function. We represent the probability density function of $T_{0,t}$ as,

$$\mathbb{P}[T_{0,t} \in d\nu] = g_{0,\tau}(\nu; x) \quad (4.18)$$

where x is the starting position of X_t at $t = 0$.

The study of occupation time distributions of random processes has attracted a great deal of attention over the years. Kac (1949), Kac and Darling (1957), and




Figure 4.1: The function $X(t) = \sin t$ and its occupation time, $T(t)$, for $a = \frac{1}{2}$.

Lamperti (1958) examined the occupation time distribution for various classes of Markov processes. More recently the occupation time distribution has been studied in a variety of physical systems such as optical imaging (Weiss and Calabrese (1996)), morphology of growing surfaces (Toroczkai et.al. (1999)), weather derivatives (Majumdar and Bray (1997)) and pricing corridor options (Fusai (2000)).

One of the many interesting results to come out of the early work of Mark Kac relates to a Markov process with a small number of states and stationary transitions. The limiting distributions for the occupation time of such a process are the Mittag-Leffler distributions. Also, for a sum of independent identically distributed random variables, the limiting distributions for the occupation time are the generalised arcsin laws. Another more widely recognised result to emerge from Kac (1949) is the derivation of the celebrated Feynman-Kac formula.

The Feynman-Kac formula arises naturally when considering the problem of

calculating the distribution of a functional, of a sequence of partial sums, of a sequence of random variables. In this section we will consider the limiting class of random walks given by the family of Brownian diffusion processes.

Consider a stochastic diffusion process X_t , given by the stochastic differential equation

$$dX_t = \alpha(X_t, t)dt + \beta(X_t, t)dW_t \quad (4.19)$$

where W_t is the Wiener process. The task before us is, given the general Wiener functional,

$$\int_0^t V(X_t)dt \quad (4.20)$$

to determine the probability density function,

$$P(x, t; \nu) = \mathbf{P}\left[\int_0^t V(X_t)dt \in \nu\right]. \quad (4.21)$$

The solution to this problem is due to Kac (1949) which states that the Laplace transform of the distribution function $P(x, t; \nu)$,

$$Q(x, t; \theta) = \int_0^\infty e^{-\theta\nu} P(x, t; \nu)d\nu \quad (4.22)$$

satisfies the partial differential equation

$$\frac{\partial Q}{\partial t} = \frac{1}{2}\beta^2(x, t)\frac{\partial^2 Q}{\partial x^2} + \alpha(x, t)\frac{\partial Q}{\partial x} - \theta V(x)Q, \quad x \in \mathbb{R}; t > 0 \quad (4.23)$$

subject to the conditions,

$$Q(x, 0) = 1$$

$$Q(x, t) \text{ bounded as } x \rightarrow \pm\infty.$$

Note that the function $Q(x, t; \theta)$ is the moment generating function of $P(x, t; \nu)$. The partial differential equation of Eq. (4.23) is commonly known as the Feynman-Kac equation. The equation takes its name from a derivation of Schrödinger's equation by theoretical physicist Richard Feynman, and the strong influence it had on the work of Kac. The Feynman-Kac equation is used in many contexts and provides a link between the theory of stochastic processes and quantum physics. In mathematical

finance it is the link between discounted expectations and the partial differential equation approach to pricing derivatives.

In the case of the occupation time we have $V(X) = \mathbf{1}(X > a)$ and the corresponding Feynman-Kac equation is,

$$\frac{\partial Q}{\partial t} = \frac{1}{2}\beta^2(x, t) \frac{\partial^2 Q}{\partial x^2} + \alpha(x, t) \frac{\partial Q}{\partial x} - \theta \mathbf{1}(x > a)Q, \quad x \in \mathbb{R}; t > 0 \quad (4.24)$$

subject to the conditions,

$$\begin{aligned} Q(x, 0) &= 1 \\ Q(x, t) &\text{ bounded as } x \rightarrow \pm\infty \\ Q(x, t), \frac{\partial Q}{\partial x} &\text{ continuous at } x = a. \end{aligned}$$

4.1.3 Deterministic activity

The first case we consider here is when the activity is specified by a deterministic function of time. We may justify such a choice by viewing the intraday patterns in Section 3.2.2 and noting that information seems to start high in the morning and decrease through until lunchtime, before picking up in the afternoon trading session. In this instance the occupation time can be calculated by evaluating the integral in Eq. (4.5) directly. For example, let the activity process be described by the function,

$$A(t) = \cos 2\pi t \quad (4.25)$$

where $0 < t < 1$ days. As seen in Fig. 4.2, the activity process provides a rough approximation to the intraday trends discovered in Fig. 3.3 of Chapter 3.

With this choice of activity the log-return of the process will correspond to that of a geometric Brownian motion with deterministic time varying parameters,

$$L_{t,\tau} = (\mu - \frac{1}{2}\sigma^2)\phi_\tau(t) + \sigma\varphi_\tau(t) [W_{t+\tau} - W_t]. \quad (4.26)$$

where

$$\varphi_\tau(t) = \int_t^{t+\tau} \mathbf{1}(\cos 2\pi s > a) ds. \quad (4.27)$$

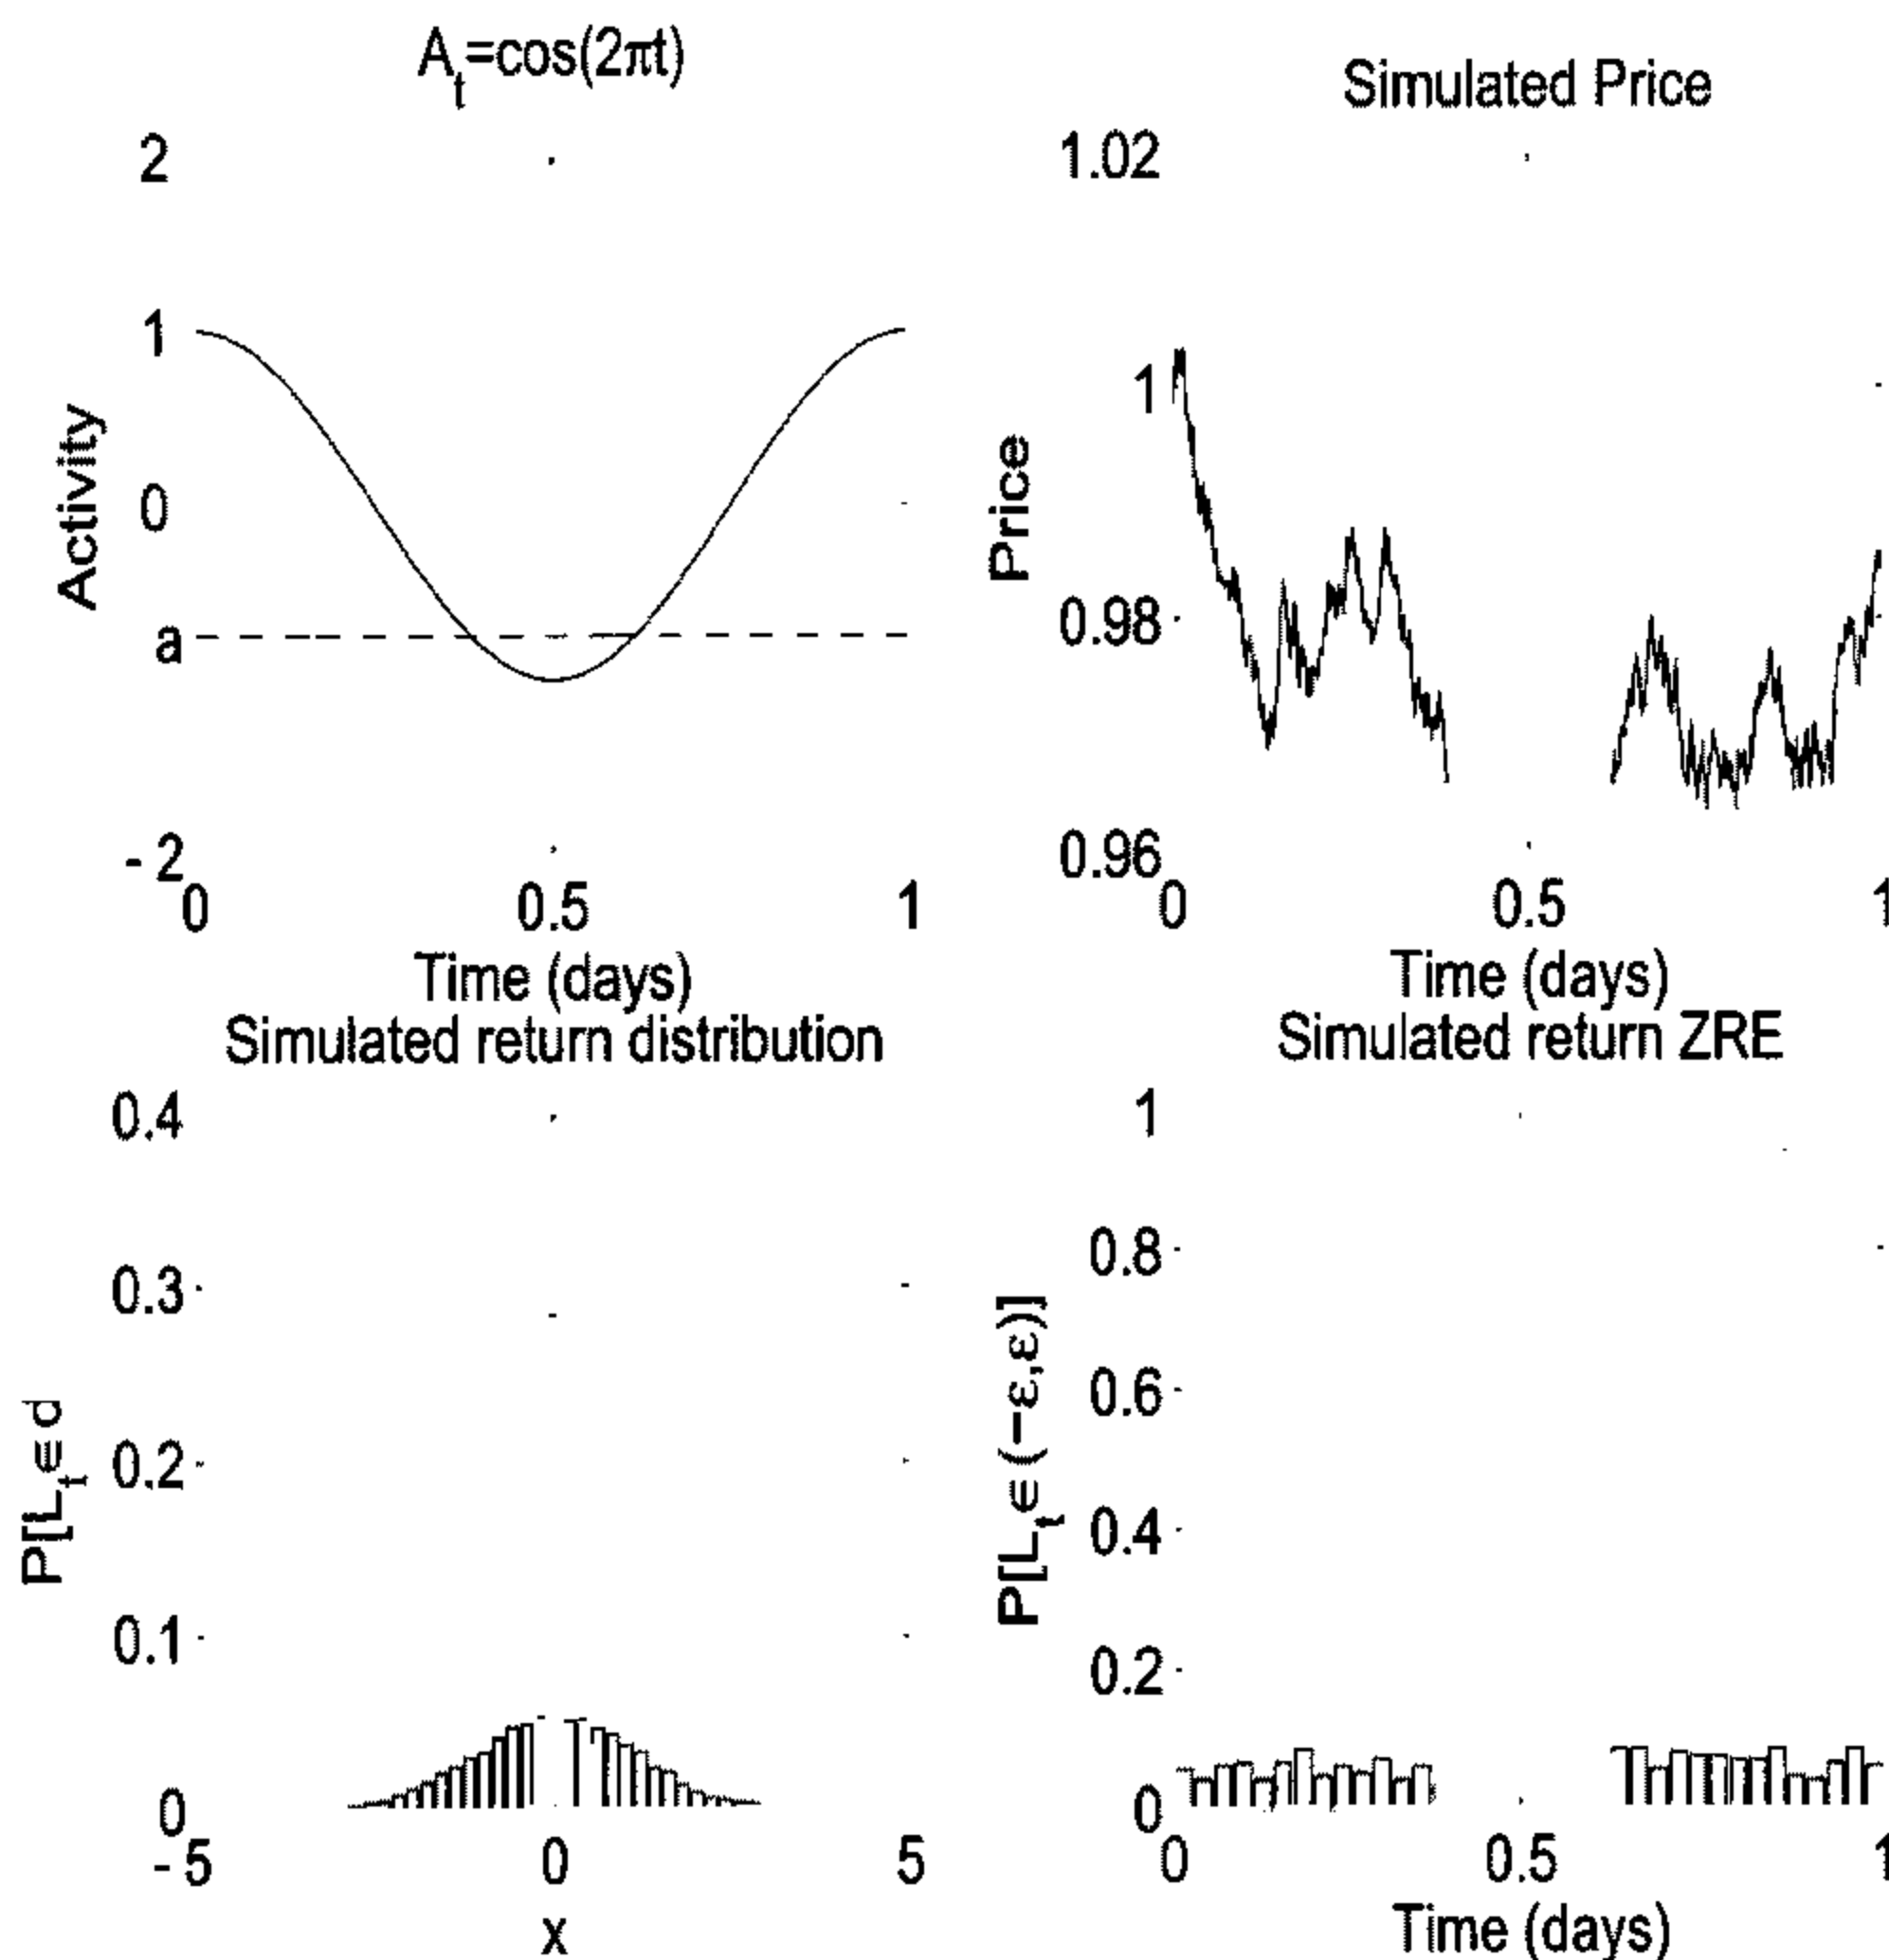


Figure 4.2: The top left figure shows the Activity as the deterministic function, $A(t) = \cos 2\pi t$. Also shown is a simulation of the price process under this choice of activity (top right), the empirical density function of the log-return (bottom left) and the intraday zero return enhancement pattern (bottom right).

Since the occupation time is a deterministic function, the integral in Eq. (4.3) disappears and the log-return will correspond to a normal distribution where the mean and variance are deterministic functions of time,

$$f_{t,\tau}(x) = \frac{1}{\sigma\sqrt{2\pi\phi_\tau(t)}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\phi_\tau(t))^2}{2\sigma^2\phi_\tau(t)}}, \quad (4.28)$$

$$\Rightarrow L_{i,\tau} \sim N\left((\mu - \frac{1}{2}\sigma^2)\phi_\tau(t), \sigma\phi_\tau(t)\right). \quad (4.29)$$

In Fig. 4.2 we show a simulation of the price process with the deterministic activity given by Eq. (4.25). The threshold level a is set such that the activity is below the threshold during the lunchtime period. We find that using a deterministic function we obtain zero return enhancement, and a corresponding intraday pattern in the stock returns. However, we see that this function produces a flat spot in the price process corresponding to time intervals where the stock will not change price with probability 1. Although this is somewhat unrealistic it represents a good first attempt at addressing the problem of zero return enhancement. To develop the model further we can replace the assumption of a deterministic activity with a randomly fluctuating activity process to account for different types of information entering the market.

4.1.4 Discrete stationary activity process

The first and simplest model for random activity considered is when the market activity is taken as a discrete stationary process. In this case we represent A_t by sequence of $n + 1$ independent identically distributed random variables, A_0, \dots, A_n . In Fig. 4.3 we see an example of such a sequence drawn from the normal distribution. Because of the discrete nature of A_t , the occupation time integral is replaced by the sum,

$$T_{0,n} = \Delta t \sum_{i=0}^n \mathbf{1}(A_i > a) \quad (4.30)$$

where there are $n + 1$ activity events in the interval $[0, \tau]$. Since A_t is a sequence of independent identically distributed random variables with distribution function $F_A(x)$, we can consider each instance of the event $\mathbf{1}(A_i > a)$, $i = 1, 2, 3, \dots, n$, as




Figure 4.3: A simulation of the activity process as a sequence of $n + 1$ discrete random variables. In this case they are drawn from the lognormal distribution. The dashed line represents the threshold level, a .

a Bernoulli trial. Thus over any interval $[t, t + \tau]$ there will be $n + 1$ successive Bernoulli trials, each with probabilities of success or failure,

$$\text{success:} \quad \mathbf{P}[\mathbf{1}(A_i > a) = 1] = p \tag{4.31}$$

$$\text{failure:} \quad \mathbf{P}[\mathbf{1}(A_i > a) = 0] = 1 - p = q. \tag{4.32}$$

Owing to the presence of the indicator function in the above equations, the probabilities p and q are obtained by calculating the probability that a random draw will be above or below a , *i.e.* ,

$$p = \mathbf{P}[\mathbf{1}(A_i > a) = 1] = \mathbf{P}[A_i > a] = 1 - F_A(a) \tag{4.33}$$

$$q = \mathbf{P}[\mathbf{1}(A_i > a) = 0] = \mathbf{P}[A_i \leq a] = F_A(a). \tag{4.34}$$

Thus the occupation time distribution for Eq. (4.30) can be specified as the distribution function for the number of successes, ν , given $n + 1$ trials each with probabilities p and q ,

$$g_{0,n}(\nu; p) = \binom{n + 1}{\nu} p^\nu q^{(n+1)-\nu}. \tag{4.35}$$

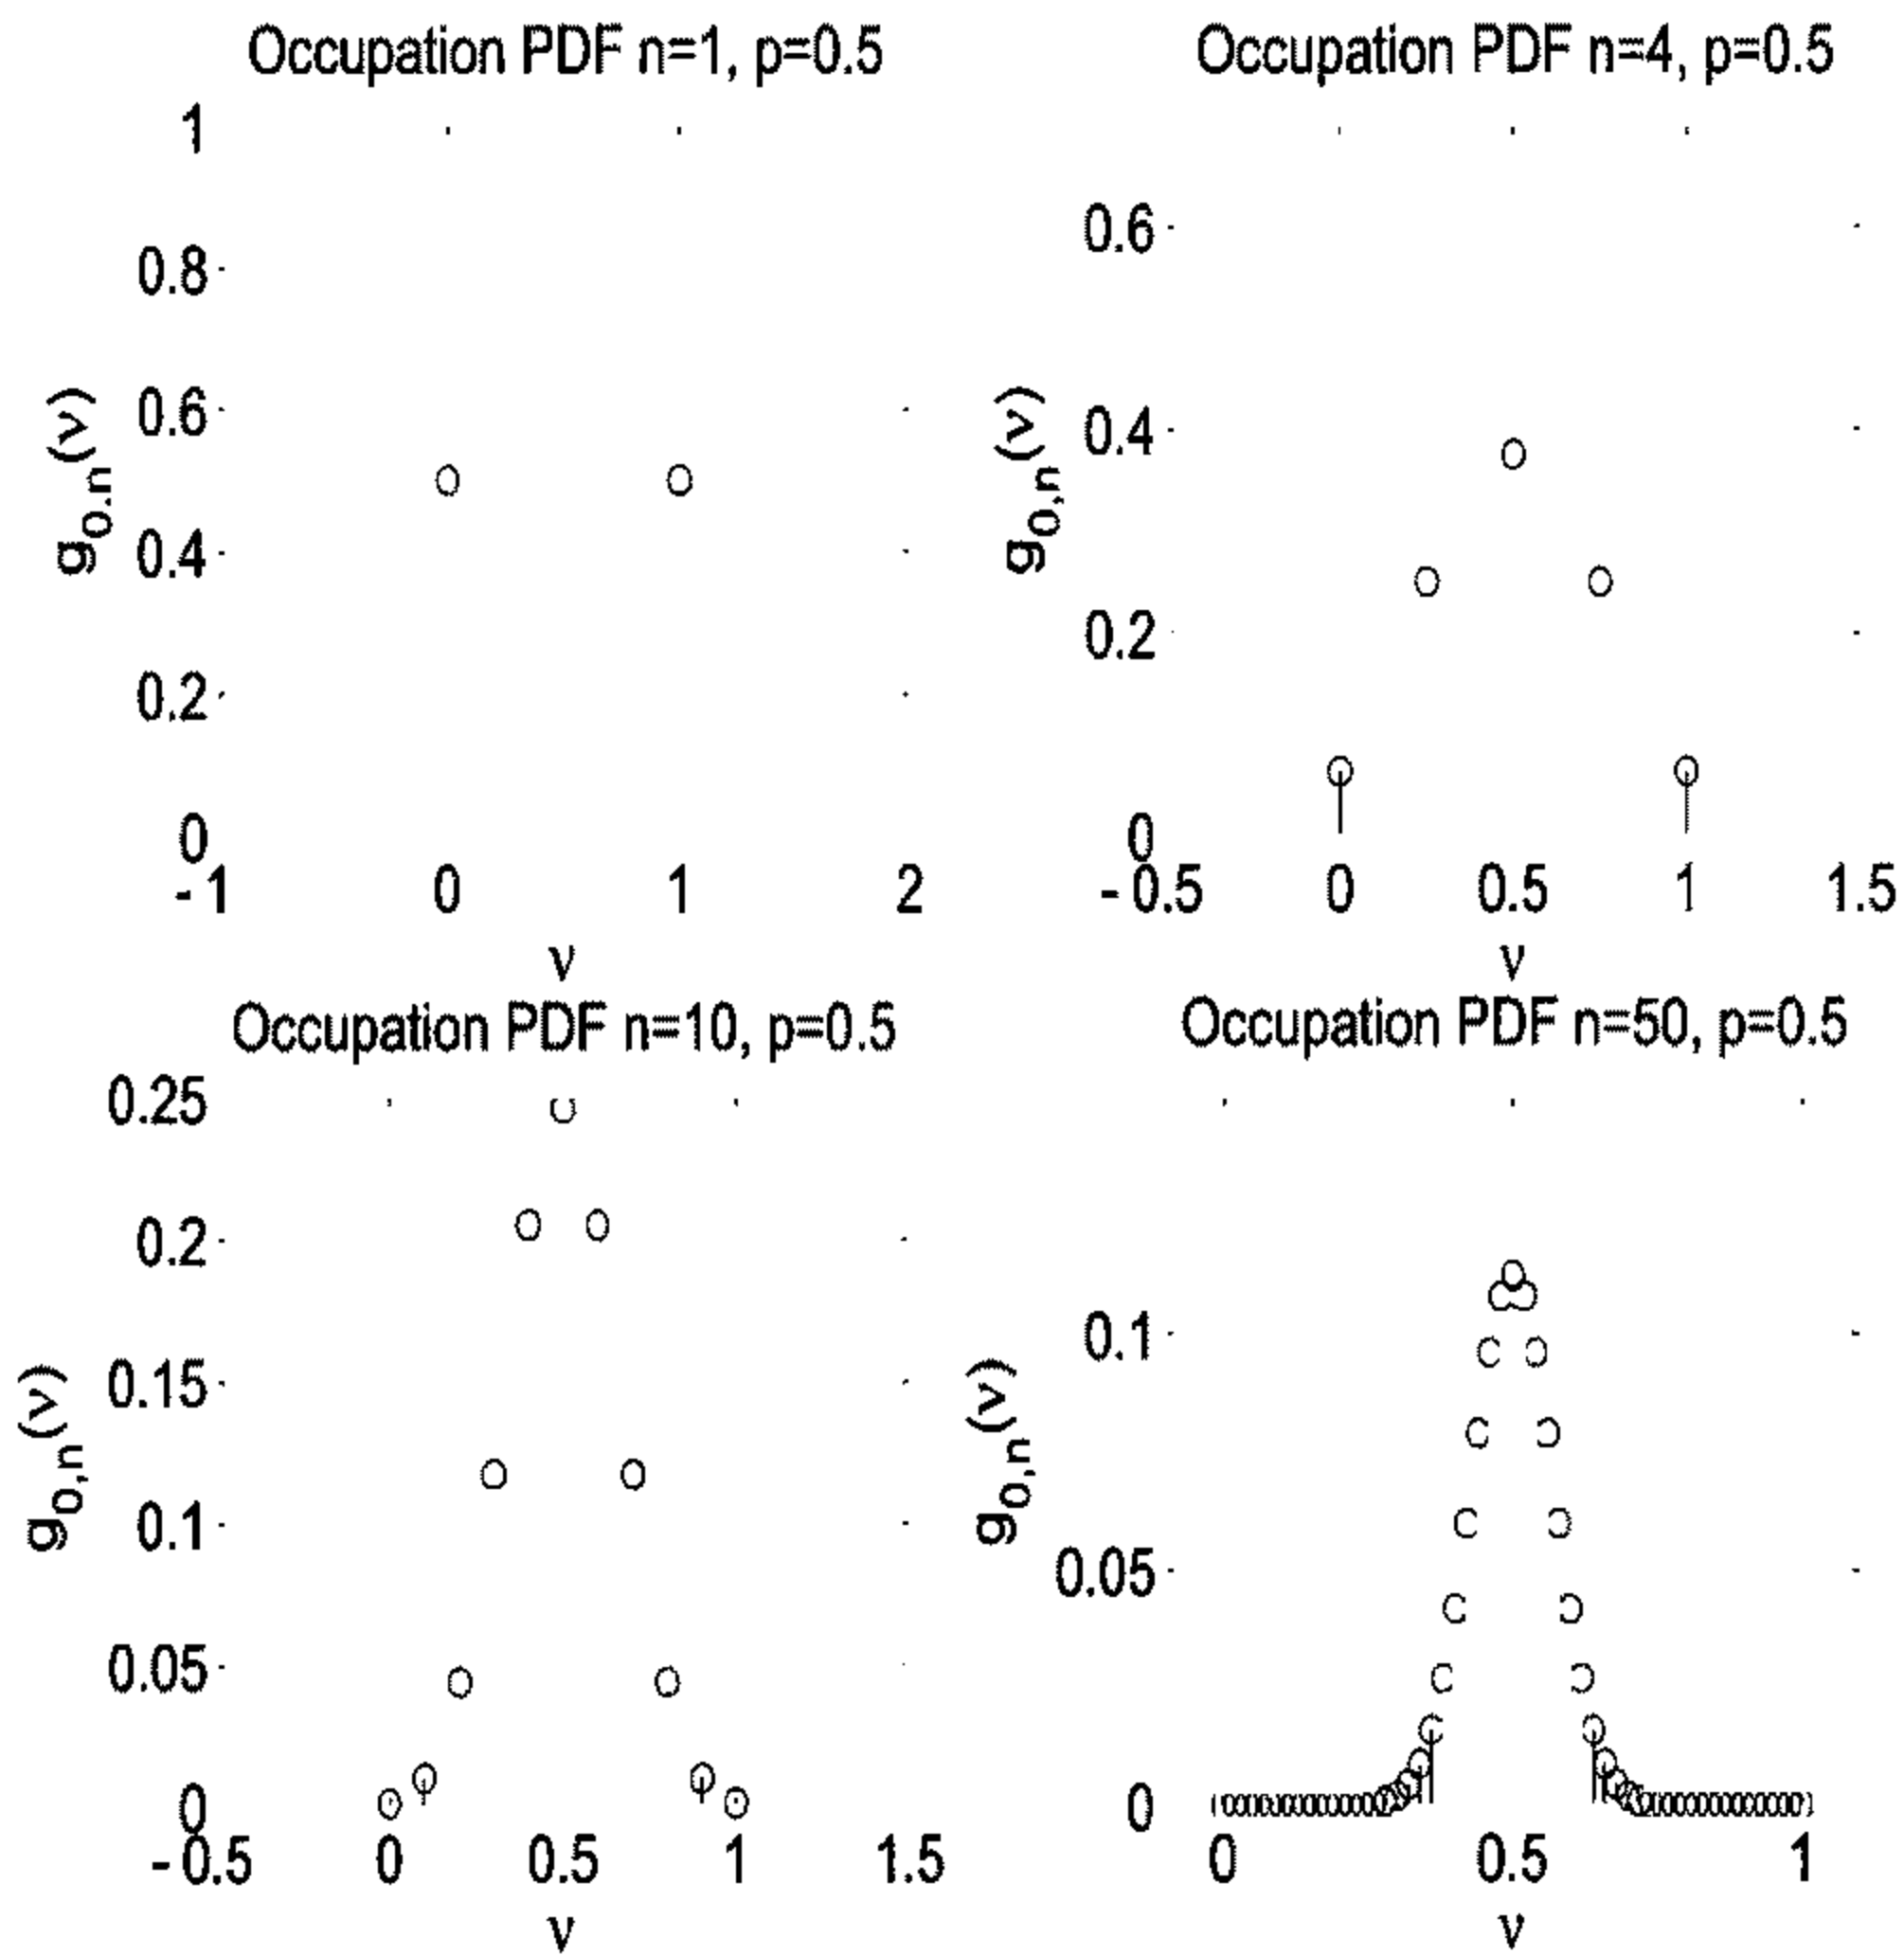


Figure 4.4: The occupation time density for a Bernoulli trial. The plots show the density functions for $n = 1, 4, 10$ and 50 activity samples per time interval.




Figure 4.5: This plot displays a sample path of the price process under the discrete activity model. We can identify periods where the price has remained constant.

For example when $n = 0$ there is only one event in the time interval τ and

$$g_{0,n}(\nu; p) = p^\nu q^{1-\nu} = \begin{cases} p & \text{for } \nu = \tau \\ q & \text{for } \nu = 0. \end{cases} \quad (4.36)$$

Fig. 4.4 displays occupation time density functions for different values of n .

We can think of this process as corresponding to the situation where there are n regularly spaced news updates per time interval. For instance, $\tau = 60$ minutes and $n = 4$ might correspond the case where information arrives in news bulletins every quarter of an hour. Depending on the impact of these bulletins, traders change their prices accordingly. This simplified model ignores the fact that news may be released at irregular times.

In Fig. 4.5 we have simulated the price process for the value $n = 50$. Fig. 4.6 displays the density functions for the log-return for, four different values of n ,

$$f_{0,n}(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_0^\tau \frac{1}{\sqrt{\nu}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\nu)^2}{2\sigma^2\nu}} \binom{n+1}{\nu} p^\nu q^{(n+1)-\nu} d\nu. \quad (4.37)$$

When n is small we see that the distribution is almost Gaussian except for a large central spike. For this choice of activity the stock price process is no longer continuous in time and the current market activity level is unrelated to past values. These

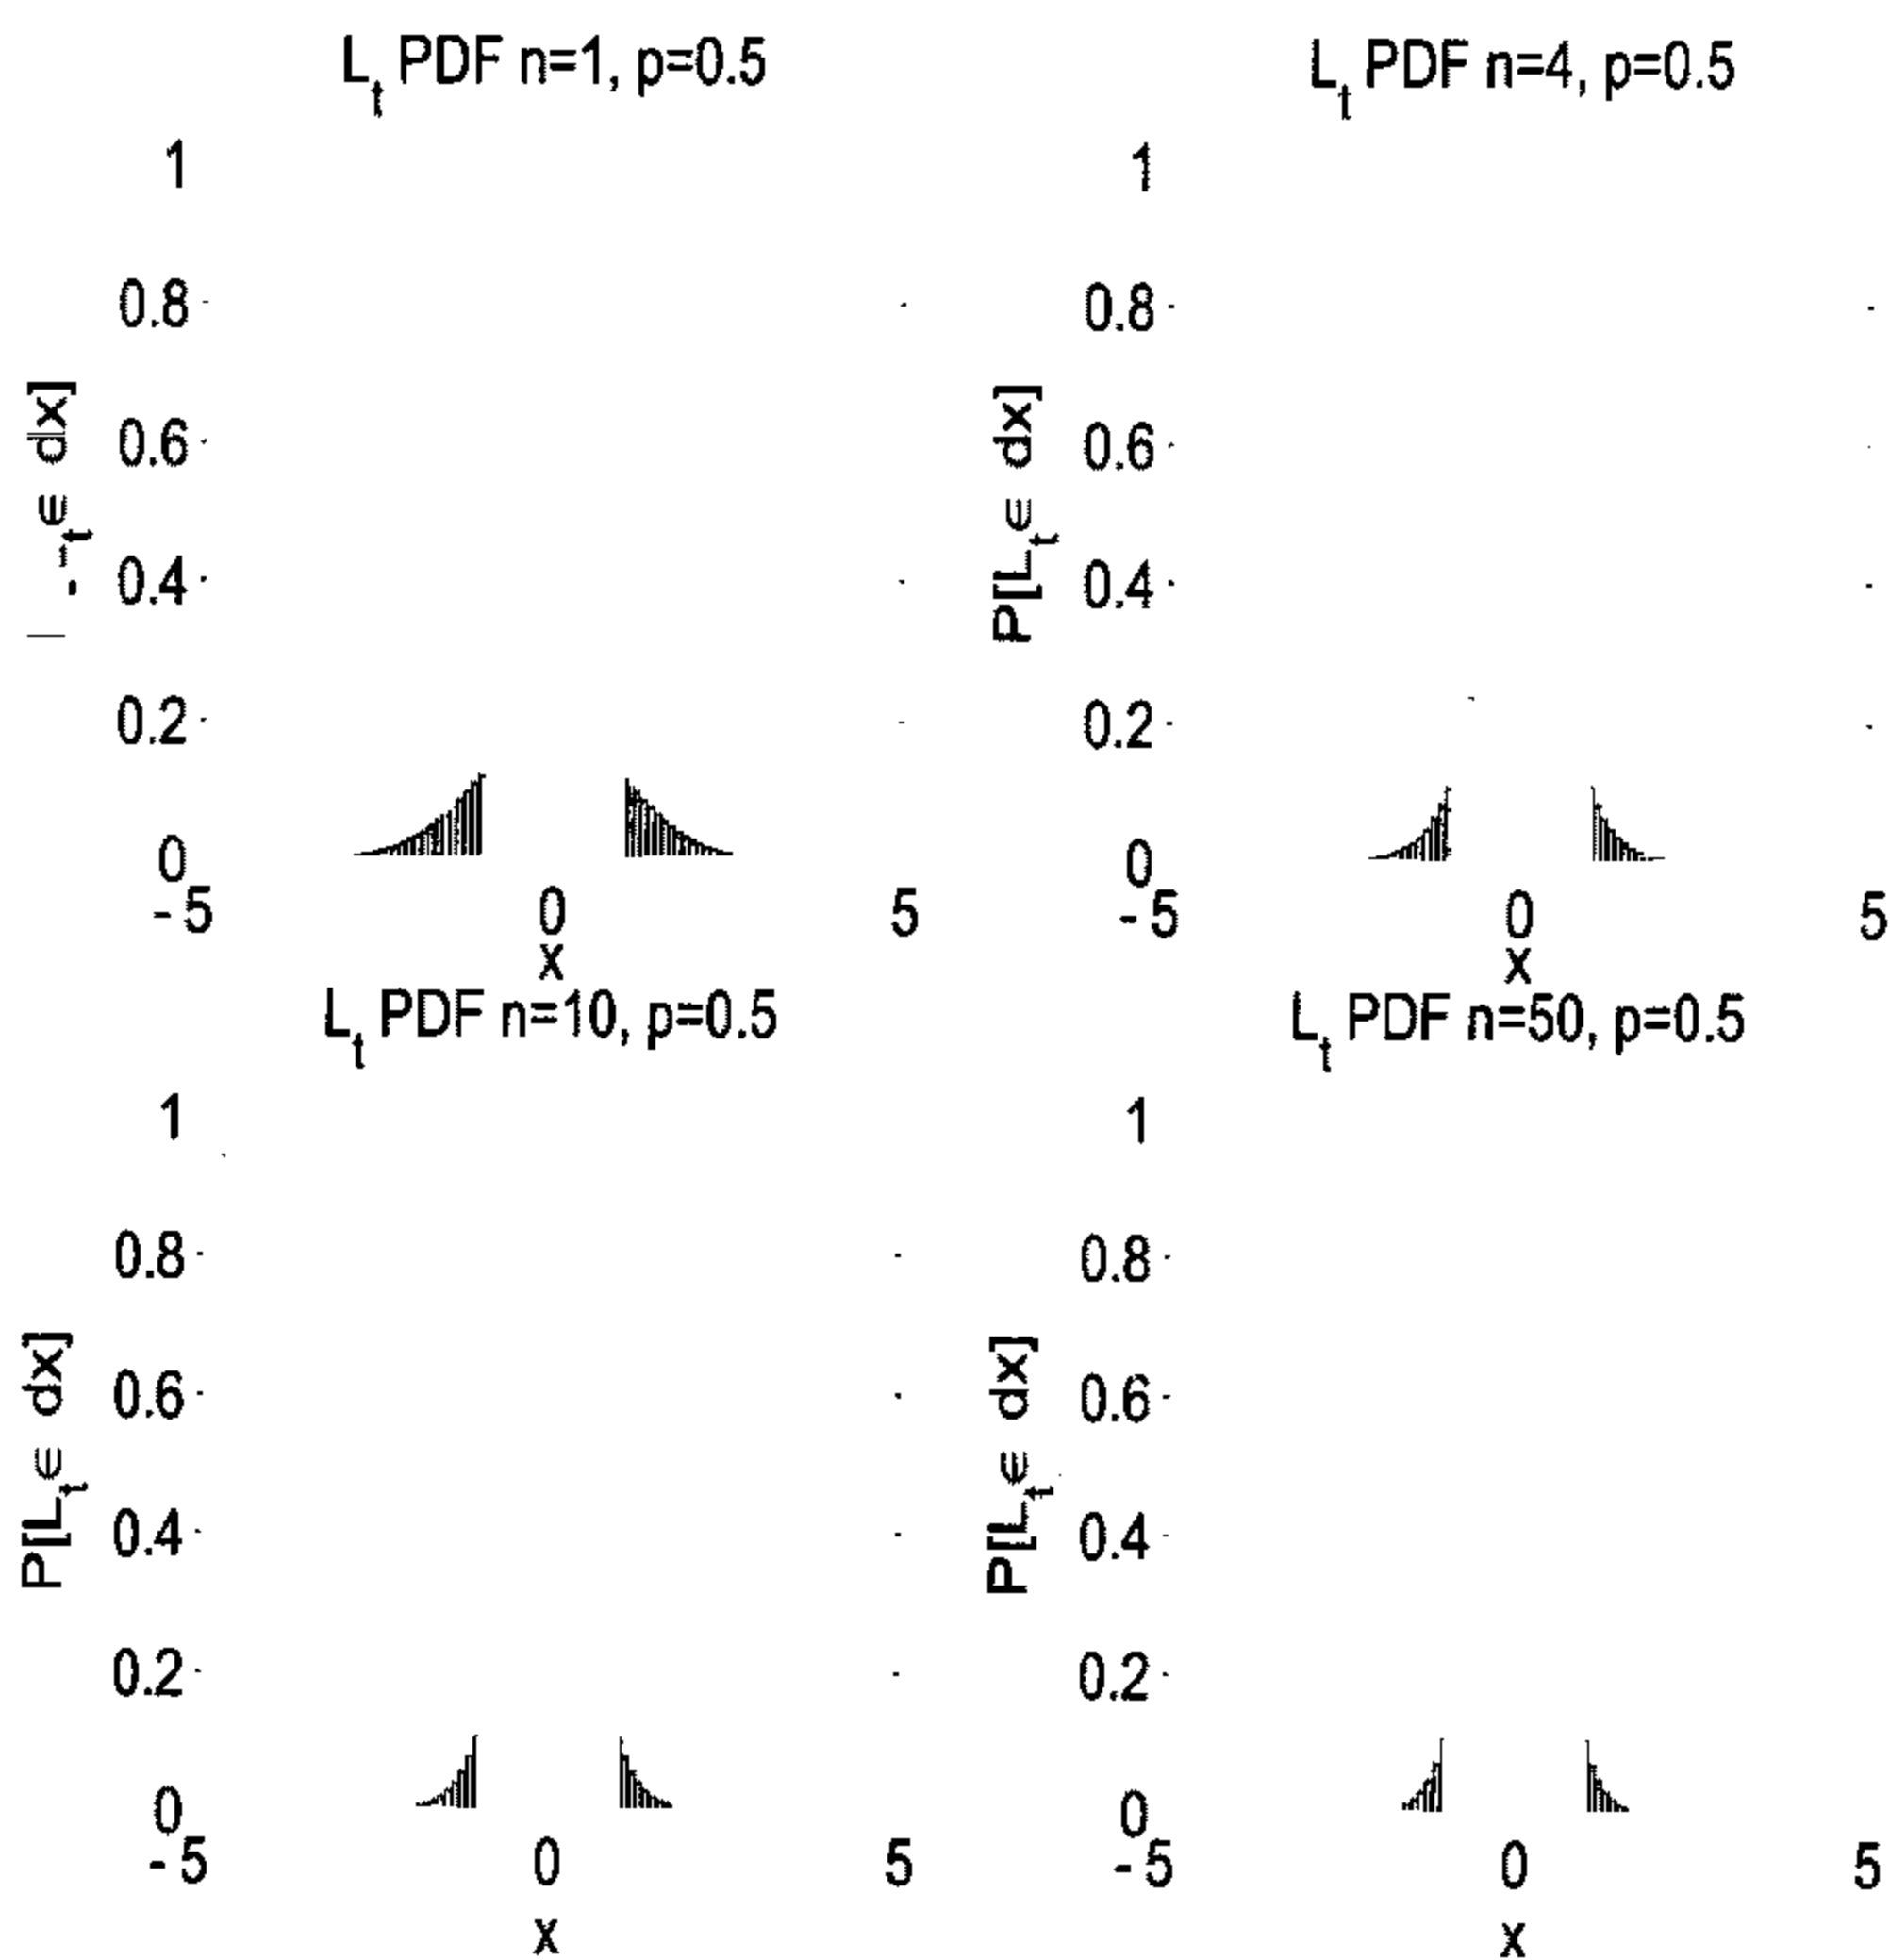


Figure 4.6: These figures display the probability density functions for the log-return with discrete activity. The plots show the density functions for $n = 1, 4, 10$ and 50 activity samples per time interval. Notice that as the number of samples per interval is increased the form of the density function changes.

deficiencies can be remedied by introducing a continuous time Markov process to model the activity.




Figure 4.7: A simulation of the market activity under the Brownian motion model.

4.1.5 Brownian activity process

In this example we let the activity process be described by the arithmetic Brownian motion,

$$A_t = m + \hat{W}_t \quad (4.38)$$

where \hat{W}_t and W_t of Eq. (4.1) are independent Wiener processes and the activity process has average level, $\mathbb{E}[A_t] = m$. This choice of activity takes into account the cumulative effect the information flow has on the market activity, with traders sentiments increasing and decreasing over time. In this model, positive information about the stock will push the activity level higher while negative information will pull the activity level lower. Fig. 4.7 provides an illustration of the Brownian activity process.

The occupation time distribution for the Wiener process has been examined in studies such as Kac (1949), Lévy (1939) and Borodin and Salminen (2002). Using a straightforward application of the Feynman-Kac equation, Eq. (4.23), with $\alpha = 0$

and $\beta = 1$, we find that the moment generating function $Q(x, t; \theta)$, satisfies the partial differential equation,

$$\frac{\partial Q}{\partial t} = \frac{1}{2} \frac{\partial^2 Q}{\partial x^2} - \theta \mathbf{1}(x > a)Q; \quad x \in \mathbb{R}, t > 0 \quad (4.39)$$

$$Q(x, 0) = 1 \quad (4.40)$$

$$Q(x, t) \text{ bounded as } x \rightarrow \pm\infty. \quad (4.41)$$

Here we have dropped the dependence on the parameter θ . Replacing x by $x - a$, this system is equivalent⁶ to

$$\frac{\partial Q}{\partial t} = \frac{1}{2} \frac{\partial^2 Q}{\partial x^2} - \theta \mathbf{1}(x > 0)Q, \quad x \in \mathbb{R}; t > 0 \quad (4.42)$$

$$Q(x, 0) = 1 \quad (4.43)$$

$$Q(x, t) \text{ bounded as } x \rightarrow \pm\infty \quad (4.44)$$

$$Q(x, t), \frac{\partial Q}{\partial x} \text{ continuous at } x = 0. \quad (4.45)$$

The Laplace transform of Eq. (4.42) in t yields the ordinary differential equation

$$s\bar{Q} - 1 = \frac{1}{2}\bar{Q}_{xx} - \theta \mathbf{1}(x > 0)\bar{Q} \quad (4.46)$$

where

$$\bar{Q} = \bar{Q}(x, s) = \mathcal{L}[Q(x, t)] \quad (4.47)$$

$$\bar{Q}_{xx} = \frac{\partial^2 \bar{Q}(x, s)}{\partial x^2} = \mathcal{L}\left[\frac{\partial^2 Q(x, s)}{\partial x^2}\right] \quad (4.48)$$

$$\bar{Q} \text{ bounded,} \quad (4.49)$$

$$\bar{Q}, \bar{Q}_x \text{ continuous at } x = 0. \quad (4.50)$$

We may now write this equation as

$$\bar{Q}_{xx} - 2s\bar{Q} = -2; \quad x < 0 \quad (4.51)$$

$$\bar{Q}_{xx} - 2(s + \theta)\bar{Q} = -2; \quad x > 0 \quad (4.52)$$

⁶This step illustrates the space homogeneous property of Brownian motion, $x + B(t) = B^x(t)$ where $B^x(t)$ is a Brownian motion started at $B(0) = x$.

which has the general solution,

$$\bar{Q}(x, s) = \begin{cases} Ae^{-\sqrt{2s}x} + \frac{1}{s}; & x < 0 \\ Be^{-\sqrt{2(s+\theta)}x} + \frac{1}{s+\theta}; & x > 0. \end{cases} \quad (4.53)$$

Using the continuity conditions at $x = 0$ the constants A and B are found to be,

$$A = \frac{1}{\sqrt{s(s+\theta)}} - \frac{1}{s} \quad (4.54)$$

$$B = \frac{1}{\sqrt{s(s+\theta)}} - \frac{1}{s+\theta}. \quad (4.55)$$

Thus the required solution to the non-homogeneous differential equation of Eq. (4.47) is given by,

$$\bar{Q}(x, s) = \begin{cases} \frac{1}{s+\theta} \left(1 - e^{-\sqrt{2(s+\theta)}x} \right) + \frac{e^{-\sqrt{2(s+\theta)}x}}{\sqrt{s(s+\theta)}}; & x > 0 \\ \frac{1}{s} \left(1 - e^{-\sqrt{2(s)}x} \right) + \frac{e^{-\sqrt{2(s)}x}}{\sqrt{s(s+\theta)}}; & x < 0. \end{cases} \quad (4.56)$$

We now need to invert the Laplace transform to obtain the moment generating function. For the case $x < 0$, we can write

$$\bar{Q}(x, s) = \frac{1}{s} - \frac{e^{\sqrt{2s}x}}{s} + \frac{e^{\sqrt{2s}x}}{\sqrt{s}} \cdot \frac{1}{\sqrt{s+\theta}}. \quad (4.57)$$

The inverse Laplace transform in s can be directly calculated using tables of standard Laplace transforms in Abramowitz and Stegun (1968),

$$Q(x, t) = \mathcal{L}_s^{-1}[\bar{Q}] \quad (4.58)$$

$$= \mathbf{1}(t > 0) - \left[1 - \text{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \right] \mathbf{1}(t > 0) \quad (4.59)$$

$$+ \frac{e^{-\frac{x^2}{2t}}}{\sqrt{\pi t}} \mathbf{1}(t > 0) \otimes \frac{e^{-\theta t}}{\sqrt{\pi t}} \mathbf{1}(t > 0) \quad (4.60)$$

$$= \text{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) + \frac{1}{\pi} \int_0^t \frac{e^{-\theta(t-s) - \frac{x^2}{2s}}}{\sqrt{s(t-s)}} ds. \quad (4.61)$$

Similarly for $x > 0$ we can write Eq. (4.56) as,

$$\bar{Q}(x, s) = \frac{1}{s+\theta} - \frac{e^{\sqrt{2(s+\theta)}x}}{s+\theta} + \frac{e^{-\sqrt{2(s+\theta)}x}}{\sqrt{s+\theta}} \cdot \frac{1}{\sqrt{s}} \quad (4.62)$$

and take the inverse Laplace transform to obtain,

$$Q(x, t) = \mathcal{L}_s^{-1}[\bar{Q}] \quad (4.63)$$

$$= e^{-\theta t} \mathbf{1}(t > 0) - e^{-\theta t} \left[1 - \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \right] \mathbf{1}(t > 0) \quad (4.64)$$

$$+ \frac{e^{-\theta t - \frac{x^2}{2t}}}{\sqrt{\pi t}} \mathbf{1}(t > 0) \otimes \frac{1}{\sqrt{\pi t}} \mathbf{1}(t > 0) \quad (4.65)$$

$$= \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) + \frac{1}{\pi} \int_0^t \frac{e^{-\theta s - \frac{x^2}{2s}}}{\sqrt{s(t-s)}} ds \quad (4.66)$$

yielding the moment generating function,

$$Q(x, s) = \begin{cases} e^{-\theta t} \operatorname{Erf} \left(\frac{x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) + \frac{1}{\pi} \int_0^t \frac{e^{-\theta s - \frac{x^2}{2s}}}{\sqrt{s(t-s)}} ds; & x > 0 \\ e^{-\theta t} \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) + \frac{1}{\pi} \int_0^t \frac{e^{-\theta(t-s) - \frac{x^2}{2s}}}{\sqrt{s(t-s)}} ds; & x < 0. \end{cases}$$

By reintroducing the dependence on θ in the moment generating function, $Q(x, t) = Q(x, t; \theta)$, it is possible to calculate the probability density function $g_{0,t}(\nu; x)$ via an inverse Laplace transform in θ , $g_{0,t}(\nu; x) = \mathcal{L}_\theta^{-1}[Q(x, t; \nu)]$. For the case $x < 0$ let $v = t - s$,

$$Q(x, t; \theta) = \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) + \frac{1}{\pi} \int_0^t \frac{e^{-\theta v - \frac{x^2}{2(t-v)}}}{\sqrt{v(t-v)}} dv \quad (4.67)$$

and the inverse Laplace transform is,

$$g_{0,t}(\nu; x) = \mathcal{L}_\theta^{-1}[Q(x, t; \nu)] \quad (4.68)$$

$$= \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) \delta(\nu) + \frac{e^{-\frac{x^2}{2(t-\nu)}}}{\pi \sqrt{v(t-v)}} \mathbf{1}(0 < \nu < t). \quad (4.69)$$

While for $x > 0$ we obtain,

$$g_{0,t}(\nu; x) = \mathcal{L}_\theta^{-1}[Q(x, t; \nu)] \quad (4.70)$$

$$= \operatorname{Erf} \left(\frac{-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) \delta(\nu - t) + \frac{e^{-\frac{x^2}{2\nu}}}{\pi \sqrt{v(t-v)}} \mathbf{1}(0 < \nu < t). \quad (4.71)$$

Finally, after substituting $x = x - a$, the occupation time density function can be

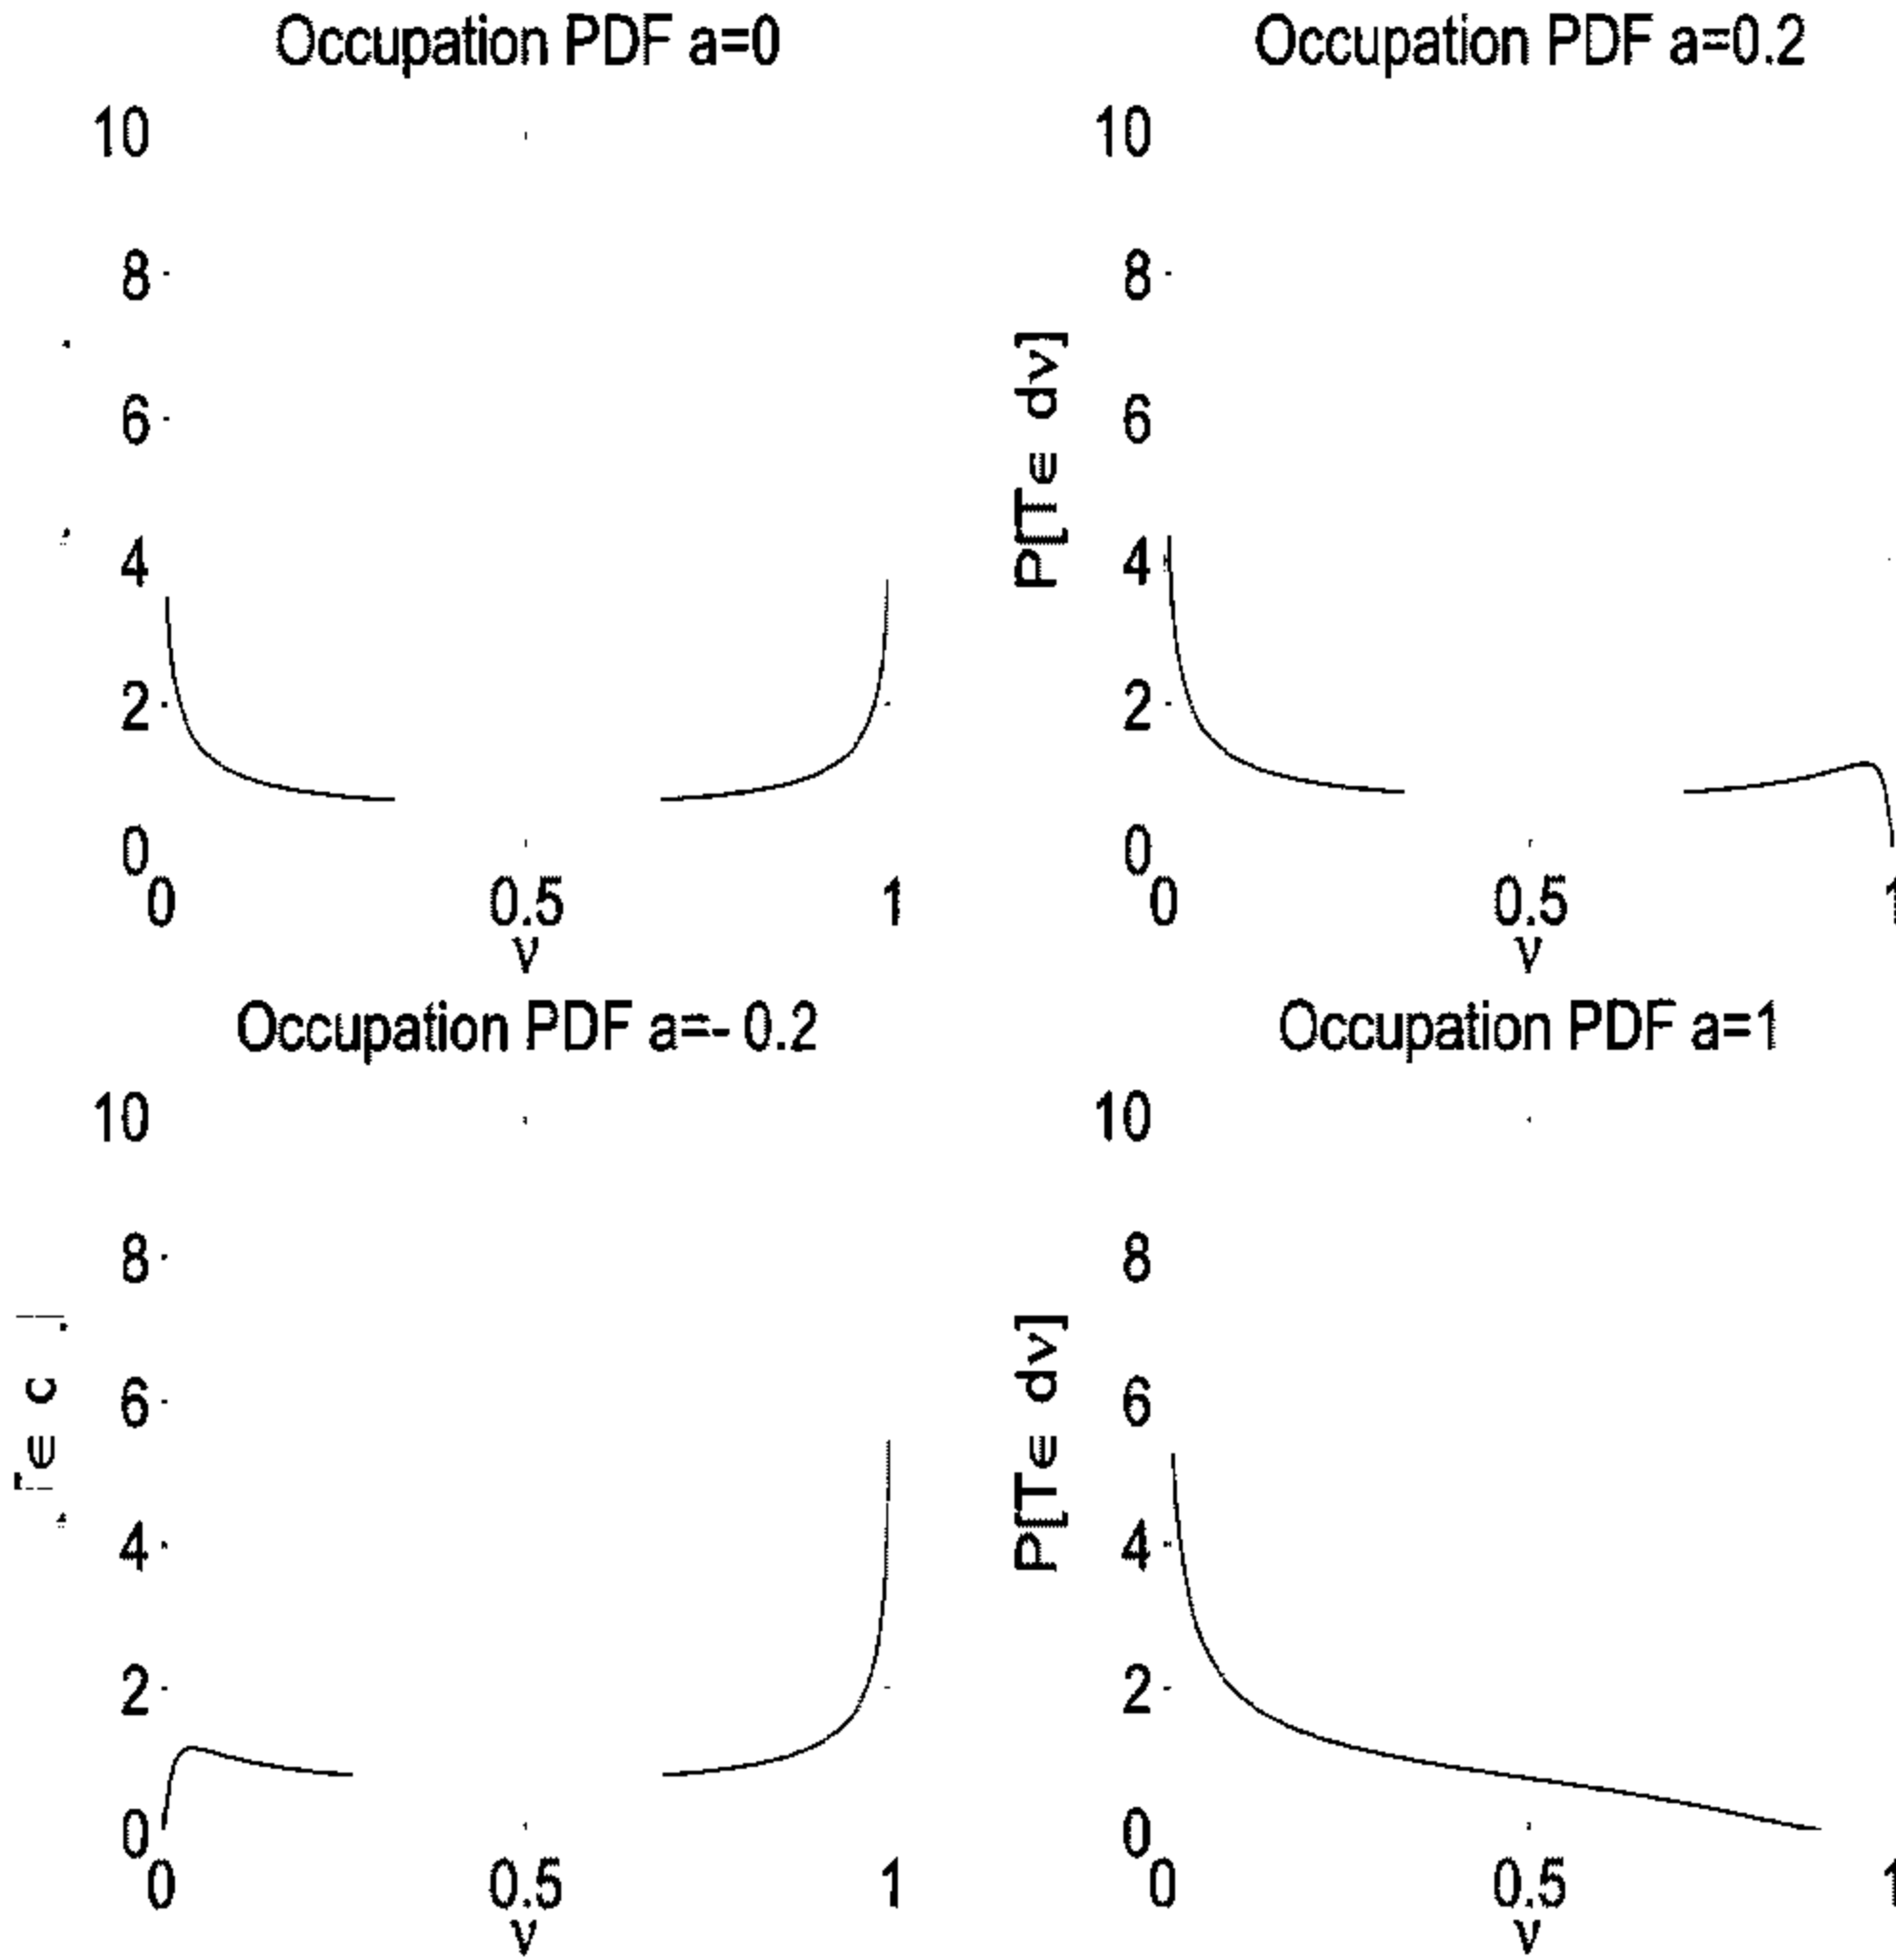


Figure 4.8: Shown here are the occupation time densities for Brownian motion with $x = 0$. At the top left we have the threshold level equal to the starting position of the Brownian motion, $a = x$. In this case the density is symmetric and indicates that the most likely outcome is that the Brownian motion will either spend all of the time above a or all of the time below a . As the starting position is changed, the nature of the density function changes as shown by the other plots in the figure.

written as,

$$g_{0,t}(\nu; x) = \begin{cases} \text{Erf} \left(\frac{a-x}{\sqrt{2t}} \right) \mathbf{1}(t > 0) \delta(\nu) + \frac{e^{-\frac{(a-x)^2}{2(\nu-t)}}}{\pi \sqrt{\nu(t-\nu)}}; & x < a \\ \text{Erf} \left(-\frac{(x-a)}{\sqrt{2t}} \right) \mathbf{1}(t > 0) \delta(\nu - t) + \frac{e^{-\frac{(x-a)^2}{2\nu}}}{\pi \sqrt{\nu(t-\nu)}}; & x > a. \end{cases} \quad (4.72)$$

Here, x is the starting position of the Wiener process \hat{W}_t and the terms $\delta(\nu)$ and $\delta(t-\nu)$ correspond to a point mass distribution and are necessary in order to satisfy the requirement,

$$\int_0^t g_{0,t}(\nu; x) d\nu = 1. \quad (4.73)$$

As shown in Fig. 4.8 the occupation time density of the Wiener process is such that if $x < a$ then the highest probability corresponds to spending the whole time interval below threshold level, *i.e.* $T_{0,\tau} = 0$. Similarly, for $x > a$, the event $T_{0,\tau} = 1$ has the highest probability. In this example where $\hat{W}_0 = 0$, Eq. (4.4) and Eq. (4.72) allow us to write down an integral representation for the probability density function of a stock's log-return over the period $[0, t]$,

$$f_{0,t}(x) = \frac{1}{\sigma \sqrt{2\pi}} \int_0^t \frac{1}{\sqrt{\nu}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\nu)^2}{2\sigma^2\nu}} g_{0,t}(\nu; x) d\nu. \quad (4.74)$$

In order to find the distribution of the log-return on $[t, t + \tau]$ we need to obtain the occupation time distribution for this time interval given that the activity process started with $\hat{W}_0 = 0$. At time t the position of the process \hat{W}_t is given by the density,

$$w(\hat{x}, t) = \frac{1}{t\sqrt{2\pi}} e^{-\frac{\hat{x}^2}{2t^2}}. \quad (4.75)$$

The required occupation time distribution on $[t, t + \tau]$ is found by conditioning Eq. (4.72) with the above density of starting positions at time t ,

$$g_{t,\tau}(\nu) = \mathbf{P} \left[\int_t^{t+\tau} \mathbf{1}(\hat{W}_s > a') ds \in d\nu \right] \quad (4.76)$$

$$= \int_{-\infty}^{\infty} w(\hat{x}, t) g_{0,\tau}(\nu; \hat{x}) d\hat{x}. \quad (4.77)$$

Thus, the density function for the log-return over the time interval $[t, t + \tau]$ is given by the double integral,

$$f_{t,\tau}(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_0^\tau \frac{1}{\sqrt{\nu}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\nu)^2}{2\sigma^2\nu}} \int_{-\infty}^\infty w(\hat{x}, t) g_{0,\tau}(\nu; \hat{x}) d\hat{x} d\nu. \quad (4.78)$$

Since we have chosen to model the activity process as a non-stationary Wiener process, we now find that a time dependence has been introduced into the log-return distribution via Eq. (4.77). This means that the log-return distribution itself is non-stationary under this particular choice of A_t .




Figure 4.9: A simulated sample path of the market activity as an Ornstein-Uhlenbeck process.

4.1.6 Ornstein-Uhlenbeck activity process

The Ornstein-Uhlenbeck process is a stationary, mean-reversion process that has been proposed as a model for volatility and interest rates by previous authors. Choosing this process to model the activity we have the stochastic differential equation,

$$dA_t = \kappa(m - A_t)dt + d\hat{W}_t \quad (4.79)$$

where m is the mean level of activity that the process reverts to and κ is the speed of reversion. Note that with a change of variables, $t = \frac{t'}{\kappa}$, $A = \frac{A'}{\sqrt{\kappa}} + m$ we obtain,

$$dA'_{t'} = -A'_{t'} dt + d\hat{W}_{t'} \quad (4.80)$$

and a new threshold level $a' = \sqrt{\kappa}(a - m)$. In this case the occupation time is,

$$T_{t',\tau} = \int_{t'}^{t'+\tau} 1(A'_s > a') ds. \quad (4.81)$$

Fig. 4.9 displays the behaviour of such an activity process. The occupation time distribution for the OU-process has been studied recently in relation to both statistical physics and weather derivatives in Majumdar and Bray (1997).

For the Ornstein-Uhlenbeck process, Eq. (4.22) has coefficients $\alpha = -\mu x$ and $\beta = 1$, hence the required Feynman-Kac PDE is,

$$\frac{\partial Q}{\partial t} = \frac{1}{2} \frac{\partial^2 Q}{\partial x^2} - \mu x \frac{\partial Q}{\partial x} - \theta \mathbf{1}(x > a) Q, \quad x \in \mathbb{R}; t > 0 \quad (4.82)$$

subject to the conditions,

$$\begin{aligned} Q(x, 0) &= 1 \\ Q(x, t) &\text{ bounded as } x \rightarrow \pm\infty \\ Q(x, t), \frac{\partial Q}{\partial x} &\text{ continuous at } x = a. \end{aligned}$$

The partial differential equation of Eq. (4.82) differs from that of Eq. (4.39) in that it contains a variable rate diffusion term, $-\mu x \frac{\partial Q}{\partial x}$. This adds another level of complexity to the problem, as replacing x by $x - a$ does not offer an simplification as it did for Brownian activity.

Taking the Laplace transform in t , we obtain the ordinary differential equation,

$$\begin{aligned} \bar{Q}_{xx} - 2\mu x \bar{Q}_x - 2(s + \theta)\bar{Q} &= -2; & x > a \\ \bar{Q}_{xx} - 2\mu x \bar{Q}_x - 2s\bar{Q} &= -2; & x < a \end{aligned} \quad (4.83)$$

where

$$\bar{Q} = \bar{Q}(x, s) = \mathcal{L}[Q(x, t)] \quad (4.84)$$

$$\bar{Q}_{xx} = \frac{\partial^2 \bar{Q}(x, s)}{\partial x^2} = \mathcal{L}\left[\frac{\partial^2 Q(x, s)}{\partial x^2}\right] \quad (4.85)$$

$$\bar{Q} \text{ bounded,} \quad (4.86)$$

$$\bar{Q}, \bar{Q}_x \text{ continuous at } x = 0. \quad (4.87)$$

Again we must solve this by matching solutions, but this time at the point $x = a$. From Kamke (1959) solutions to an equation of the form $y'' - 2axy' - 2(s+k)y = 0$ are given in terms of the Hermite functions. In the case of Eq. (4.83) the general solution is found to be

$$\bar{Q}(x, s) = \begin{cases} \frac{1}{\theta+s} + C_1 \mathbf{H}[-\frac{\theta-s}{\mu}, \sqrt{\mu}x]; & x > a \\ \frac{1}{s} + C_2 \mathbf{H}[-\frac{s}{\mu}, \sqrt{\mu}x]; & x < a \end{cases} \quad (4.88)$$

where $\mathbf{H}[a, x]$ are the Hermite functions and C_1, C_2 are constants. Using the relationship,

$$\frac{d}{dx} \mathbf{H}[n, ax] = 2na \mathbf{H}[n-1, ax] \quad (4.89)$$

we match \bar{Q} and \bar{Q}_x at $x = a$ to obtain the coefficients,

$$C_1 = \frac{\theta}{\theta+s} \frac{\mathbf{H}[-1-\frac{s}{\mu}, \sqrt{\mu}a]}{s \mathbf{H}[-1-\frac{s}{\mu}, \sqrt{\mu}a] \mathbf{H}[-\frac{\theta-s}{\mu}, \sqrt{\mu}a] - (\theta+s) \mathbf{H}[-\frac{s}{\mu}, \sqrt{\mu}a] \mathbf{H}[-1+\frac{-\theta-s}{\mu}, \sqrt{\mu}a]}; \quad x > a \quad (4.90)$$

$$C_2 = \frac{\theta}{s} \frac{\mathbf{H}[-1+\frac{-\theta-s}{\mu}, \sqrt{\mu}a]}{s \mathbf{H}[-1-\frac{s}{\mu}, \sqrt{\mu}a] \mathbf{H}[-\frac{\theta-s}{\mu}, \sqrt{\mu}a] - (\theta+s) \mathbf{H}[-\frac{s}{\mu}, \sqrt{\mu}a] \mathbf{H}[-1+\frac{-\theta-s}{\mu}, \sqrt{\mu}a]}; \quad x < a.$$

When $a = 0$, this represents the situation where the threshold level is equal to the mean reversion level of the process. In this case, matching \bar{Q} and \bar{Q}_x gives the coefficients in terms of gamma functions, hence the solution for $a = 0$ is

$$\bar{Q}(x, s) = \begin{cases} \frac{1}{\theta+s} + \frac{-\theta}{\theta+s} \frac{\Gamma(\frac{\theta+s}{\mu}+1) \Gamma(\frac{1}{2}+\frac{s}{2\mu}) \mathbf{H}[-\frac{\theta-s}{\mu}, \sqrt{\mu}x]}{(s+\theta) \Gamma(1+\frac{s}{2\mu}) \Gamma(\frac{1}{2}+\frac{\theta+s}{2\mu}) - s \Gamma(1+\frac{\theta+s}{2\mu}) \Gamma(\frac{1}{2}+\frac{s}{2\mu})}; & x > 0 \\ \frac{1}{s} + \frac{-\theta}{s} \frac{\Gamma(\frac{s}{\mu}+1) \Gamma(\frac{1}{2}+\frac{\theta+s}{2\mu}) \mathbf{H}[-\frac{s}{\mu}, \sqrt{\mu}x]}{(s+\theta) \Gamma(1+\frac{s}{2\mu}) \Gamma(\frac{1}{2}+\frac{\theta+s}{2\mu}) - s \Gamma(1+\frac{\theta+s}{2\mu}) \Gamma(\frac{1}{2}+\frac{s}{2\mu})}; & x < 0. \end{cases} \quad (4.91)$$

In deriving this solution we used the Gamma function identity,

$$\Gamma[z+1] \Gamma[z+\frac{1}{2}] = \sqrt{\pi} \frac{\Gamma[2z+1]}{2^{2z}}. \quad (4.92)$$

In order to calculate the occupation time density function we must perform a double inverse Laplace transform. However, due to the complexity of the solution it is not

even possible to obtain a closed form for the moment generating function $Q(x, t; \theta)$ for either the general case Eq. (4.90), or for the $a = 0$ case Eq. (4.91), let alone transform this solution to obtain the function $g_{0,t}(\nu; x)$. Although it may be possible to attempt this calculation numerically, such a procedure would prove troublesome as the numerical inversion of Laplace transforms can be a very difficult computation.

Instead, we adopt a different approach to calculate the occupation time distribution. Though the Ornstein-Uhlenbeck process is a diffusion process, it exhibits mean reversion behaviour. This means that given a sufficiently long time interval, the transient effects will disappear and the process will become stationary. Although we are concerned mainly with the long term behaviour of the process, it is this transient behaviour that adds complication to calculating the occupation time distribution. To avoid this, we implement a method due to Majumdar and Bray (1997) which allows us to calculate the occupation time distribution over a sufficiently large time interval τ .

For the OU process,

$$dX_t = -X_t dt + dW_t \quad (4.93)$$

we define the variable u for an arbitrary function $V(X)$,

$$u = \frac{1}{\tau} \int_0^\tau V(X_s) ds \quad (4.94)$$

such that $T_{0,\tau} = u\tau$ when $V(X) = \mathbf{1}(X > a)$. The Laplace transform of the density function $g_{0,\tau}(\nu)$ can be written in terms of the moment generating function of $T_{0,\tau}$,

$$\mathcal{L}[g_{0,\tau}(\nu)] = \mathbb{E}_{T_{0,\tau}}[e^{-u\tau\theta}] = \frac{Z(\theta)}{Z(0)} \quad (4.95)$$

where $Z(\theta)$ is given by the path integral

$$Z(\theta) = \int DX_\tau e^{-\frac{1}{2} \int_0^\tau ds [\dot{X}^2 + 2\dot{X}X + X^2 + 2\theta V(X_s)]}. \quad (4.96)$$

Since we are interested in large τ , choosing periodic boundary conditions, $X(T) = X(0)$ will not affect the results in the large τ limit. This has the benefit of removing the perfect derivative term $2X\dot{X}$ so that,

$$Z(\theta) = \int DX_\tau e^{-\frac{1}{2} \int_0^\tau ds [\dot{X}^2 + X^2 + 2\theta V(X_s)]}. \quad (4.97)$$

This integral gives the partition function of a quantum particle with Hamiltonian, $\mathcal{H} = p^2/2 + X^2/2 + sV(X)$. For $\tau \rightarrow \infty$ the ground state dominates to give

$$\mathbb{E}_{T_0, \tau}[e^{-u\tau\theta}] = e^{-\tau[E(\theta) - E(0)]} \quad (4.98)$$

where $E(\theta)$ is the ground state energy eigenvalue for the Schrödinger equation,

$$-\frac{1}{2} \frac{d^2\psi}{dX^2} + \left[\frac{X^2}{2} + \theta V(X) \right] \psi = E(\theta)\psi. \quad (4.99)$$

For $\theta = 0$ the problem is that of the simple harmonic oscillator with $E(0) = \frac{1}{2}$. Using the method of steepest descents to invert the Laplace transform in Eq. (4.95) Majumdar and Bray (1997) obtain the general form,

$$g_{0, \tau}(\nu) \sim e^{-\Phi(u)\tau} \quad (4.100)$$

where

$$\Phi(u) = \max_{\theta} [E(\theta) - \frac{1}{2} - u\theta] \quad (4.101)$$

is called the large deviation function.

To calculate the occupation time distribution we choose the function $V(X) = \mathbf{1}(X > a)$. Then the differential equation to be solved is,

$$\psi'' - X^2\psi + 2(E(\theta) - \theta)\psi = 0; \quad X > a \quad (4.102)$$

$$\psi'' - X^2\psi + 2E(\theta)\psi = 0; \quad X < a \quad (4.103)$$

subject to the conditions,

$$\psi(X) \quad \text{bounded as } X \rightarrow \pm\infty \quad (4.104)$$

$$\psi(X), \psi'(X) \quad \text{continuous at } X = a. \quad (4.105)$$

The solution to equations of this form is given by Kamke (1959) in terms of the confluent hypergeometric function and can be expressed as parabolic cylinder functions, $D_n(x)$ (see Abramowitz and Stegun (1968)),

$$\psi(X) = \begin{cases} \psi_+ = C_1 D_{n_+}(\sqrt{2}X); & X > a \\ \psi_- = C_2 D_{n_-}(-\sqrt{2}X); & X < a \end{cases} \quad (4.106)$$

with

$$n_+ = -\frac{1}{2} + E(\theta) - s \quad (4.107)$$

$$n_- = -\frac{1}{2} + E(\theta). \quad (4.108)$$

The coefficients C_1 and C_2 are determined by solving $\psi_+(a) = \psi_-(a)$ and $\psi'_+(a) = \psi'_-(a)$. This gives the equation,

$$\frac{(-1 + 2E)D_{-1+n_-}(-\sqrt{2}a)}{D_{n_-}(-\sqrt{2}a)} = \frac{(-1 + 2E - s)D_{-1+n_+}(\sqrt{2}a)}{D_{n_+}(\sqrt{2}a)}. \quad (4.109)$$

The function $E(\theta)$ can be determined from the above eigenvalue equation. Numerical solution of this equation requires careful choice of initial guess to avoid singularities. Once $E(\theta)$ is found we use Eq. (4.100) and Eq. (4.101) to calculate the required large deviation and density functions. Appendix E contains the program that solves Eq. (4.109) for $E(\theta)$. Fig. 4.10 displays $\Phi(u)$ calculated for two different values of a . Also shown are the corresponding density functions compared with results obtained by simulation.

The occupation time distribution of the OU-process

$$g_{t,\tau}(\nu) \sim e^{\Phi(u)\tau} \quad (4.110)$$

differs greatly from that of the Wiener process. In the large τ limit, the OU-process is stationary and so the occupation time distribution is independent of the starting position A_0 . Also the functional form of Eq. (4.110) is such that the probabilities for $T_{0,\tau} = 0$ or $T_{0,\tau} = \tau$ are very small. The stationary property of the OU-process means that for large τ , the distribution of $T_{t,\tau}$ is invariant under translation, *i.e.*

$$\mathbf{P} \int_0^\tau \mathbf{1}(A_s > a) ds \sim \mathbf{P} \int_t^{t+\tau} \mathbf{1}(A_s > a) ds \quad (4.111)$$

$$\Rightarrow g_{0,\tau}(\nu) = g_{t,\tau}(\nu). \quad (4.112)$$

Hence the probability density function of the log-return for large τ ,

$$f_{t,\tau}(x) \sim \frac{1}{\sigma\sqrt{2\pi}} \int_0^\tau \frac{1}{\sqrt{\nu}} e^{-\frac{(x - (\mu - \frac{1}{2}\sigma^2)\nu)^2}{2\sigma^2\nu}} g_{0,\tau}(\nu) d\nu \quad (4.113)$$

is stationary.

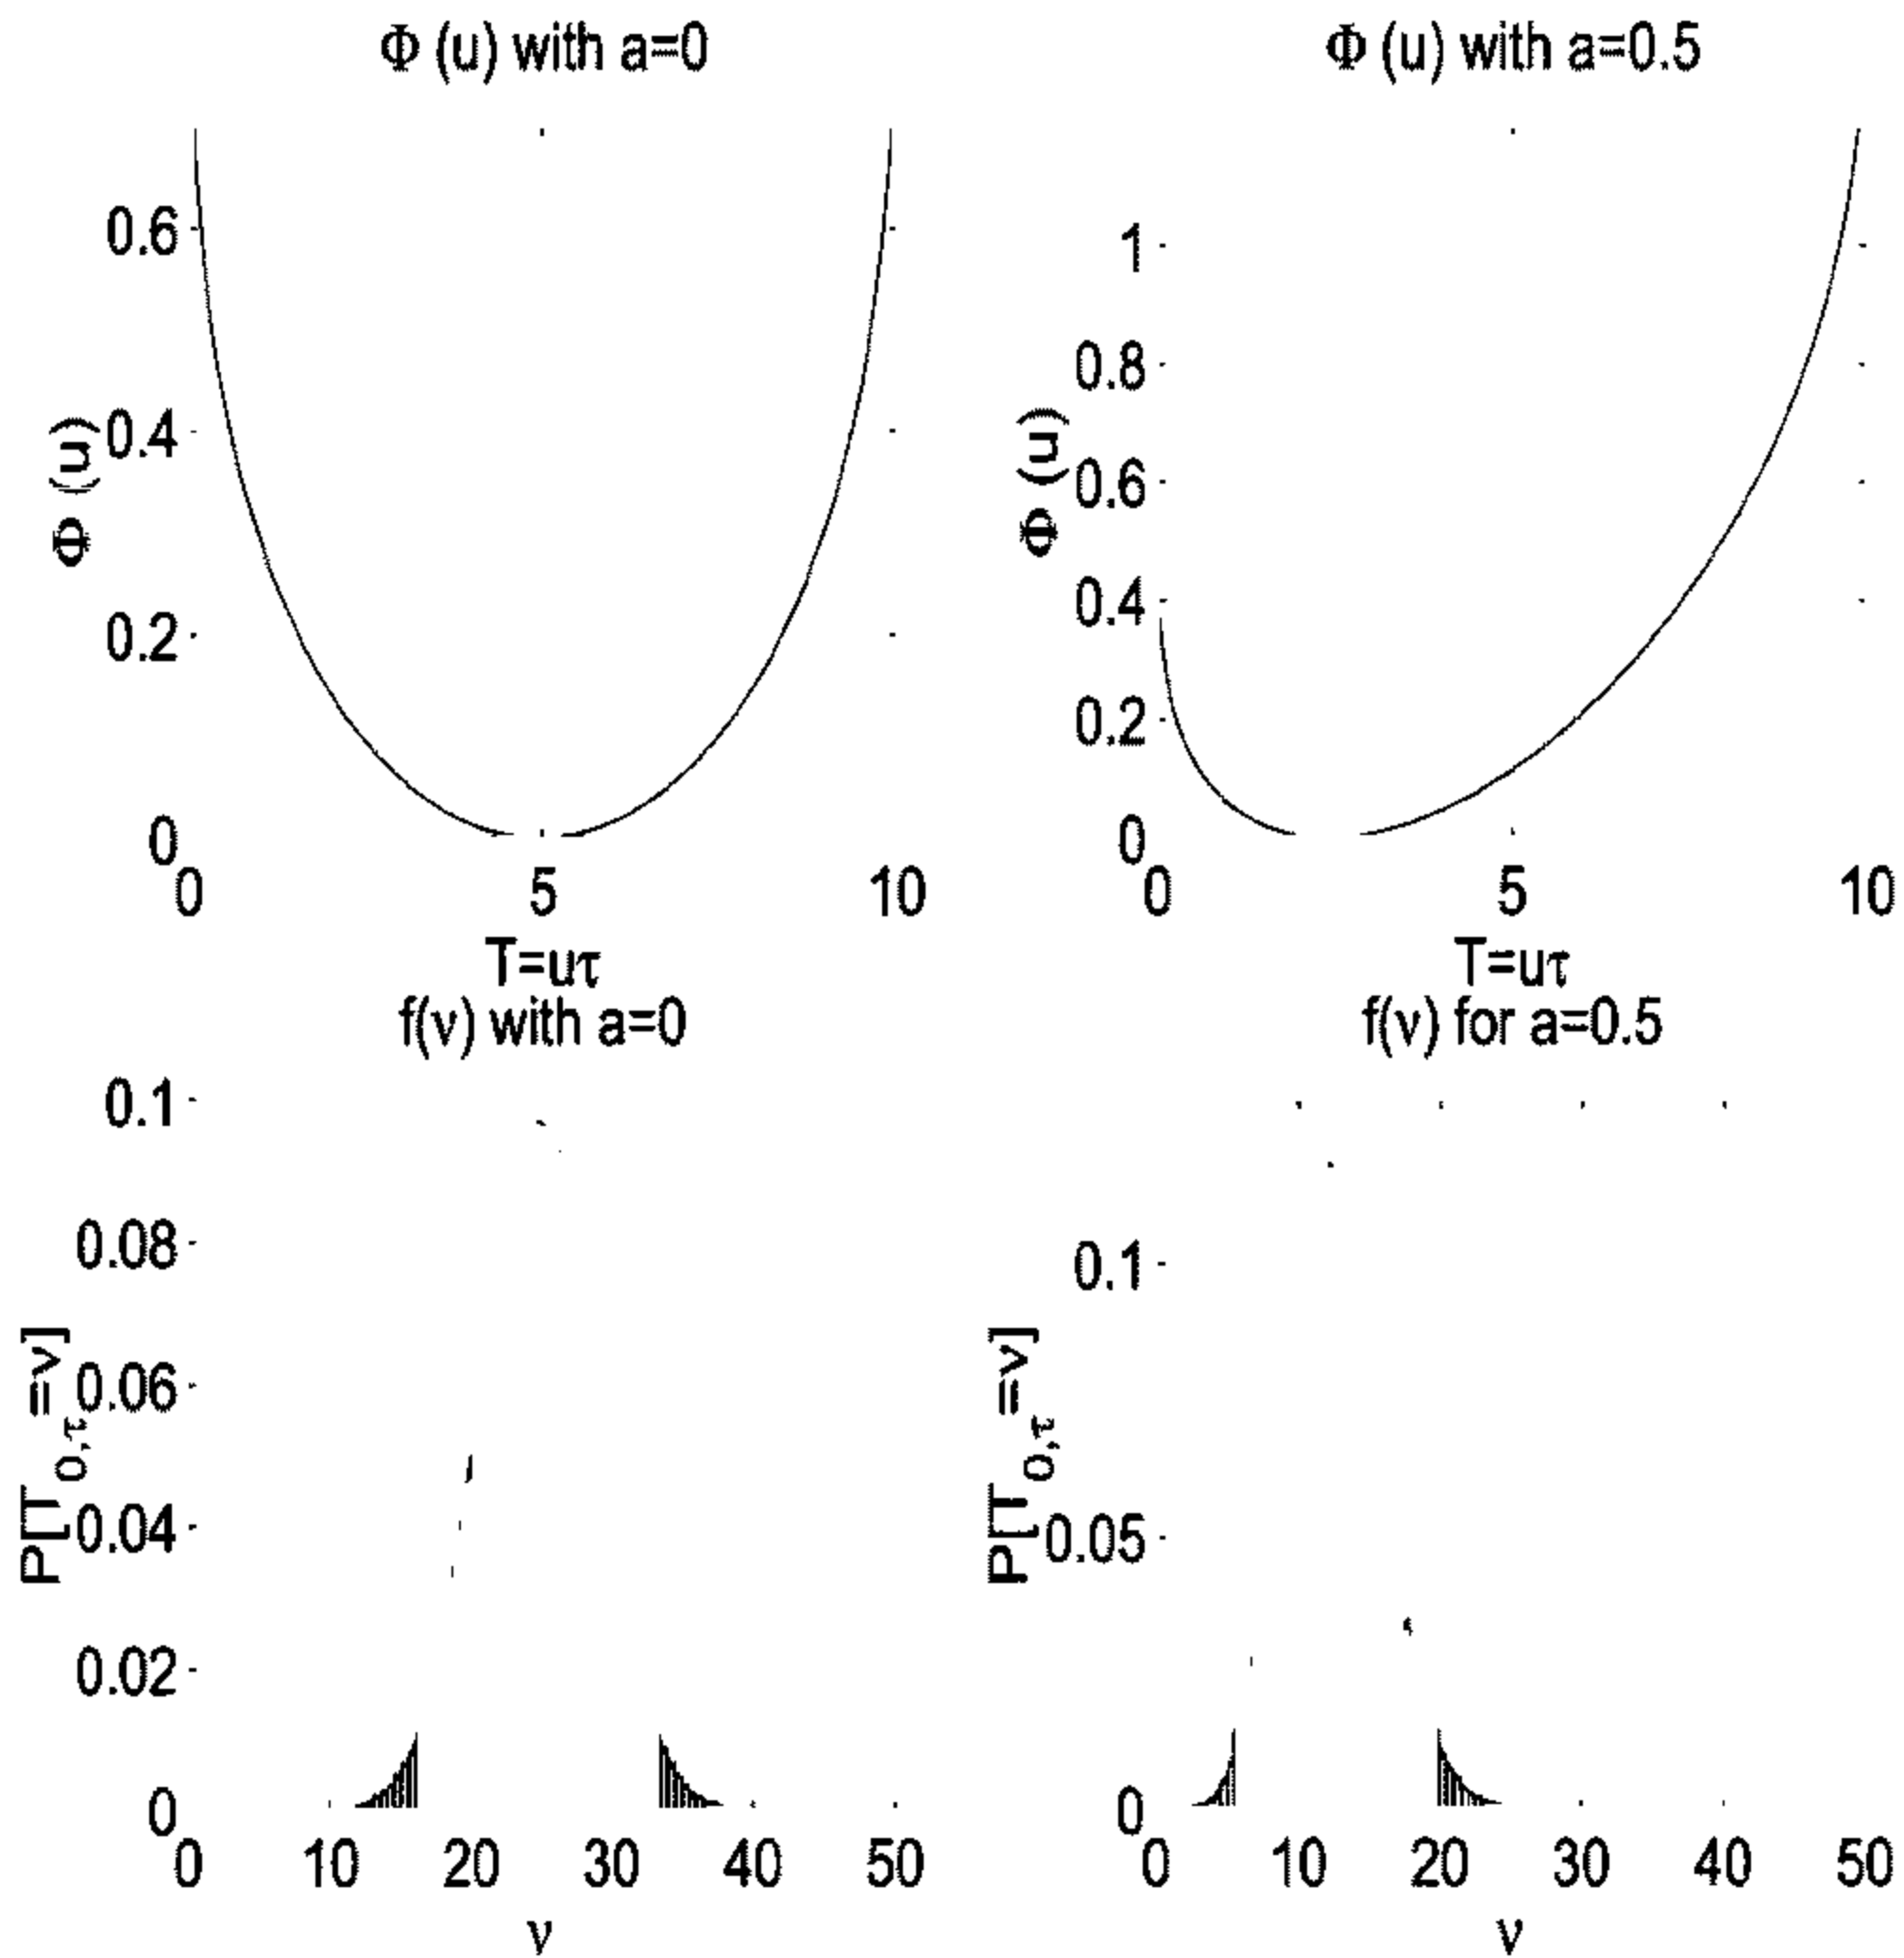


Figure 4.10: The top two plots show the large deviation function for $a = 0$ and $a = 0.5$. The bottom plots show the large τ occupation time density for $a = 0$ and $a = 0.5$. The bar plots in the bottom figures represent the simulated occupation density.

4.1.7 Results of fitting to ASX data

In this section we fit the model to data using two activity processes from the previous section. With these two example processes we approximate the behaviour of the activity process for small and large time scales and show that threshold Brownian motion is able to describe the behaviour of the log-return over small time interval τ_S , and large time interval τ_L . Working in the business time scale constructed with the master volatility measure, we take the activity in each interval to be equally distributed. We assume the market activity to behave like a stationary process for large τ_L , while over small τ_S the activity process exhibits local non-stationary behaviour. That this might be the case follows from the following reasoning. In the information age, more and more information is available to traders. Where once market moving information may have been reported only sparsely, now with the advent of the 24hr news cycle, such information is arriving more often. As this information begins to arrive with greater frequency, traders initially trade more on these events. However as this level of information becomes more commonplace, traders build up a resistance to this information and begin to take such events in their stride. With the increase in information flow comes a corresponding increase in the resistance to trade on this information. Thus the trading activity will behave in a non-stationary manner over short time scales, but over a longer time scale will revert to some average level.

Interval τ_S

Over small intervals τ_S we expect to see numerous instances of zero return enhancement. This means that there is a large probability that the activity process will be below the threshold level for the whole time interval τ_S . For each small time interval $[t, t + \tau_S]$, we approximate the activity by a Wiener process with random starting position \hat{x}_t . The distribution of starting positions is taken to be independent of interval starting time, thus for each time interval $[t, t + \tau_S]$ the activity will be equal in distribution and when viewed on a larger time scale the activity process will be stationary. Fig. 4.11 shows the threshold model fitted to the CBA 10 minute log-return. Here, the behaviour of the Wiener process and its occupation time dis-




Figure 4.11: This figure displays the results of fitting the TBM model to normalised CBA log-return data in business time. Shown on the left is the fitted empirical density for the 10 minute log-return where we have modelled the small interval activity as a Wiener process. On the right is the fitted density for the 1 hour log-return where we have modelled the large interval activity as an OU-process.

tribution have allowed for the specific modelling of the large number of zero returns. The values of Pearson’s correlation coefficient in Table C.1 indicate that our model is in excellent agreement with the data.

Interval τ_L

For large intervals, the probability of finding zero return enhancement decreases, as it is highly unlikely that a stock will trade with constant price for a large length of time. Correspondingly we expect the activity process to have a low probability of spending the whole time interval completely above or below the threshold level. We choose to model the large τ activity behaviour with the Ornstein-Uhlenbeck process noting that over sufficiently small time scales the OU-process behaves like a

Wiener process⁷. Fitting the threshold model to the CBA 1 hour log-return we do not see the zero return enhancement that was so apparent in the previous example. However, zero returns are still occurring on small scales in $[t, t + \tau_L]$ and their effect on the large scale behaviour of the model is apparent in Fig. 4.11. The log-return density of the model exhibits a marked deviation from the Gaussian behaviour of GBM, displaying fatter tails and a much sharper central peak. This result illustrates the strong effects that small scale market events can have on large scale log-return behaviour. Fig. 4.11 represents a remarkably good approximation to the log-return distribution. Examination of the tail behaviour in Figure 4.12 shows that TBM describes the data well. While the model appears to slightly underestimate the 10 minute log-return, it provides an accurate representation of the tail behaviour of the 1 hr log-return distribution.




Figure 4.12: This figure displays tail behaviour of the TBM model compared to the CBA log-return distribution.

⁷The Ornstein-Uhlenbeck process could be used to model both the small and large time intervals if it were possible to calculate the occupation time density of the transient part of the OU-process.

4.2 Montroll-Weiss continuous time random walk

A recent contribution to the modelling of high frequency financial markets has been made by Scalas et.al. (2000). This Italian research group presents a general phenomenological theory for the tick-by-tick dynamics of financial markets. There has been much interest in this model and the conceptual approach it offers. The model they implement centers around the idea that in financial markets not only the price can be modelled as a random variable, but the waiting time between two consecutive transactions is stochastic. It is an idea similar to this that underlies the theory and methods of subordination. However, in much of the subordination literature the time variable is taken as an economic measure of trading time.

In mathematical finance the price dynamics of financial markets are represented as random walks. In order to make use of many desirable mathematical features, these random walks are studied in continuous time, a method pioneered by Bachelier. There are in fact other ways to embed a random walk in continuous time. Scalas et.al. (2000) use the approach of Montroll and Weiss (1967) that produces the so-called continuous time random walk (CTRW). Under this process, the time interval between successive steps is itself a random variable. The CTRW was developed by mathematicians working on probability theory for the diffusion process on lattices and is described in detail by Hughes (1995). In economics, the CTRW model goes back to Hilfer (1994) in the context of sale monitoring. Scalas et.al. (2000) in their model demonstrate that the Lévy-scaling form follows as a particular case. Their approach is able to take into account the non-Markovian and non-local character of financial time series.

In this section of Chapter 4 we describe the continuous time random walk and the corresponding application to financial modelling. We then investigate the model empirically using high frequency ASX data, to estimate waiting times and compare these with the theoretic predictions.

4.2.1 CTRW theory

A continuous time random process, such as Brownian motion, is usually obtained by taking a joint limit of $\Delta \rightarrow 0$ and $\tau \rightarrow 0$ of a random walk on a lattice, where Δ is the lattice spacing and τ is the time between successive steps. Such processes are governed by the Lévy-stable densities. This approach, while leading to the intricacies of advanced probability theory, destroys the lattice structure and the conceptual simplicity of the random walk. An alternative approach that yields a continuous time random walk is due to Montroll and Weiss (1967). This model is a random walk where the time between steps, T_j , are independent identically distributed random variables with distribution function $\psi(t)$. We regard the step as taking place instantaneously with the walker resting in between steps.

Under this model the random walk is characterised by a joint density function $\phi(l, \tau)$, of jump size l and waiting time τ . The two probability distributions,

$$\lambda(l) := \int_0^{\infty} \phi(l, \tau) d\tau \quad (4.114)$$

$$\psi(\tau) := \int_{-\infty}^{\infty} \phi(l, \tau) dl \quad (4.115)$$

are known as the ‘jump pdf’ and the ‘waiting-time pdf’. For the Montroll-Weiss CTRW, the probability of jump size is taken to be independent of waiting time. Thus the joint density function admits the factorisation,

$$\phi(l, \tau) = \lambda(l)\psi(\tau). \quad (4.116)$$

In the case where this factorisation is not possible we have the more general Scher-Lax model (Scher and Lax (1973)). The Montroll-Weiss model is used because it allows for easy deduction from the discrete time random walk.

We obtain the properties of the CTRW by first defining the discrete time random walk. Let $\lambda(l)$ denote the transition probability for taking a jump of length l . Then we define the probability of being at position l after n jumps as,

$$\mathbf{P}[X_n = l] = P_n(l) \quad (4.117)$$

with initial condition

$$P_0(l) = \delta_{l,0}. \quad (4.118)$$

We now introduce the generating function,

$$P(l; \xi) = \sum_{n=0}^{\infty} P_n(l) \xi^n \quad |\xi| < 1. \quad (4.119)$$

The evolution of the discrete time random walk can be deduced from the probability $P_{n+1}(l)$ that a walker departing from the origin is found at site l after $n + 1$ jumps. If the walker is at position l' after n jumps, then arrival at the required site on the next jump has probability $\lambda(l - l')$. Summing over all possible positions l' yields,

$$P_{n+1}(l) = \sum_{l'} \lambda(l - l') P_n(l'). \quad (4.120)$$

This difference equation describes the evolution of the walk. Eq. (4.120) can be solved by first taking the discrete Fourier transform (DFT). Write,

$$\tilde{P}_n(k) = \sum_l e^{ilk} P_n(l) \quad (4.121)$$

and define the 'structure function' of the walk as,

$$\gamma(k) = \sum_l e^{ilk} \lambda(l). \quad (4.122)$$

Now, the DFT of Eq. (4.120) is given by

$$\tilde{P}_{n+1}(k) = \gamma(k) \tilde{P}_n(k) \quad (4.123)$$

with initial condition

$$\tilde{P}_0(k) = 1. \quad (4.124)$$

Thus the solution to Eq. (4.123) is found to be

$$\tilde{P}_n(k) = \gamma(k)^n \quad (4.125)$$

and inverting the DFT yields,

$$P_n(l) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ilk} \gamma(k)^n dk. \quad (4.126)$$

Thus for a given structure function it is possible to evaluate $P_n(l)$. In order to remove the n dependence in Eq. (4.126) we can look instead at the generating function. Multiplying both sides of Eq. (4.126) by ξ^n and summing over all n yields,

$$\sum_{n=0}^{\infty} P_n(l)\xi^n = \sum_{n=0}^{\infty} \xi^n \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ilk} \gamma(k)^n dk \quad (4.127)$$

and via Eq. (4.120) we obtain,

$$P(l; \xi) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ilk} \sum_{n=0}^{\infty} \gamma(k)^n \xi^n dk \quad (4.128)$$

$$= \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{e^{-ilk}}{1 - \gamma(k)\xi} dk \quad (4.129)$$

From this integral we may obtain the statistics of the random walk.

In order to obtain the CTRW we now make use of the Montroll-Weiss theorem. The Montroll-Weiss theorem provides a way of deducing the properties of a CTRW given the generating function of the corresponding discrete time random walk. Specifically, the theorem states that if

$$P(l; \xi) = \sum_n P_n(l)\xi^n \quad (4.130)$$

is the generating function for a DTRW then for the corresponding CTRW with waiting time density, $\psi(t)$, the probability of occupying position x at time t , $p(x, t)$, has Laplace transform,

$$\hat{p}(x, u) = \frac{1 - \hat{\psi}(u)}{u} P(x; \hat{\psi}(u)) \quad (4.131)$$

where $\hat{\psi}(u)$ is the Laplace transform of the waiting time pdf, $\psi(t)$.

Consider the case of DTRW where $P(x; \xi)$ is given by Eq. (4.129), then using Eq. (4.131) and Eq. (4.129) the corresponding CTRW has,

$$\hat{p}(x, u) = \frac{1 - \hat{\psi}(u)}{2\pi u} \frac{1}{2\pi} \int_{-\pi}^{\pi} \frac{e^{-ilk}}{1 - \gamma(k)\hat{\psi}(u)} dk \quad (4.132)$$

Taking the DFT of the above equation yields,

$$\tilde{\hat{p}}(k, u) = \frac{1 - \hat{\psi}(u)}{u [1 - \gamma(k)\hat{\psi}(u)]} \quad (4.133)$$

This equation is known as the Montroll-Weiss equation. For inversion of the Fourier and Laplace transforms, we make use of the convolution theorem to obtain the evolution equation for $p(x, t)$,

$$p(x, t) = \delta(x)\Psi(t) + \int_0^t \psi(t-t') \int_{-\infty}^{\infty} \lambda(x-x')p(x', t')dx' dt' \quad (4.134)$$

where $\Psi(t) = 1 - \int_0^t \psi(t')dt'$ is the survival probability. The survival probability represents the probability that the diffusing quantity does not change value during a time interval t after a jump. Eq. (4.134) is commonly called the master equation of the CTRW. The probability of being at position x at time t is thus the combined probability that the walker is in state x' at time t' and waits $t - t'$ units of time before making a jumping of length $x - x'$, evaluated for all possible values of x' and t' .

Scalas et.al. (2000) and Mainardi et.al. (2000) present an alternative form of the general master equation for the CTRW. This form can be interpreted as a Fokker-Plank type evolution equation. To derive this form, we first rewrite Eq. (4.133) as

$$\tilde{\Phi}(s) [s\hat{p}(k, s) - 1] = \hat{p}(k, s) \quad (4.135)$$

where

$$\tilde{\Phi}(s) = \frac{1 - \hat{\psi}(s)}{s\hat{\psi}(s)} = \frac{\hat{\Psi}(s)}{\hat{\psi}(s)} = \frac{\hat{\Psi}(s)}{1 - s\hat{\Psi}(s)} \quad (4.136)$$

The master equation in Eq. (4.134) then becomes,

$$\int_0^t \Phi(t-t') \frac{\partial}{\partial t'} p(x, t') dt' = -p(x, t) + \int_{-\infty}^{\infty} \lambda(x-x')p(x', t) dx' \quad (4.137)$$

with $\Phi(t)$ defined such that

$$\Psi(t) = \int_0^t \Phi(t-t')\psi(t')dt'. \quad (4.138)$$

In this new form of master equation, the rate at which transitions occur is determined not only by the current state, but by the entire past history of the process through the memory kernel $\Phi(t)$. As a consequence of this, the CTRW is in general, a non-Markovian process.

The CTRW will be memoryless and hence Markovian if the memory kernel degenerates to a delta function (multiplied by a positive constant), *i.e.* $\Phi = \mu\delta(t)$. In this case we find

$$\hat{\psi}(s) = \frac{\mu}{1+s} \quad (4.139)$$

$$\hat{\Psi}(s) = \frac{1}{1+s} \quad (4.140)$$

and thus the waiting time pdf, $\psi(t)$, and the survival function, $\Psi(t)$, are the exponential densities,

$$\psi(t) = \mu e^{-\mu t} \quad (4.141)$$

$$\Psi(t) = e^{-\mu t}. \quad (4.142)$$

Under these conditions the new master equation reduces to

$$\mu \frac{\partial}{\partial t} p(x, t) = -p(x, t) + \int_{-\infty}^{\infty} \lambda(x-x') p(x', t) dx', \quad p(x, 0) = \delta(x). \quad (4.143)$$

This equation is the Kolmogorov-Feller equation, the most general form of master equation for a Markovian CTRW.

4.2.2 CTRW in finance

As we have seen in previous chapters, financial data is dependent on past values. In fact, financial time series display memory behaviour that decays as a power law of time. We saw in the last section that if the memory kernel is a delta function we recover the Markovian CTRW. In order to include the effect of long memory into the CTRW, Mainardi et.al. (2000) consider a non-Markovian process characterised by a memory kernel that exhibits a power law decay. They choose the function,

$$\Phi(t) = \frac{-t^\beta}{\Gamma(1-\beta)} \quad (4.144)$$

with $t \geq 0$ and $0 < \beta < 1$. For the limiting case $\beta = 1$, the memory kernel reduces to $\Phi(t) = \delta(t)$. Taking the Laplace transform of $\Phi(t)$,

$$\tilde{\Phi}(s) = \frac{1}{s^{1-\beta}} \quad (4.145)$$

we can substitute into Eq. (4.135) to obtain,

$$s^\beta \hat{p}(k, s) - s^{\beta-1} = [\tilde{\lambda}(k) - 1] \hat{p}(k, s); \quad 0 < \beta < 1. \quad (4.146)$$

In this case the master equation becomes,

$$\frac{\partial^\beta}{\partial t^\beta} p(x, t) = -p(x, t) + \int_{-\infty}^{\infty} \lambda(x - x') p(x', t') dx'; \quad p(x, 0) = \delta(x) \quad (4.147)$$

where $\frac{\partial^\beta}{\partial t^\beta}$ is the Caputo fractional derivative of order β , defined in the Laplace domain by the relation,

$$\mathcal{L} \left[\frac{\partial^\beta}{\partial t^\beta} f(t) \right] = s^\beta \tilde{f}(s) - s^{\beta-1} f(0^+) \quad 0 < \beta < 1. \quad (4.148)$$

This fractional operator has been used extensively in the field of linear viscoelasticity by Caputo and Mainardi (1971) and Buchen and Mainardi (1975) to model dissipation effects. Thus, Eq. (4.147) can be considered as the time-fractional generalisation of the Kolmogorov-Feller equation.

The choice of $\Phi(t)$ in Eq. (4.144) implies a particular form for the waiting-time pdf and the survival function. Using Eq. (4.136) we find that,

$$\tilde{\psi}(s) = \frac{1}{1 + s^\beta} \quad (4.149)$$

$$\tilde{\Psi}(s) = \frac{s^{\beta-1}}{1 + s^\beta} \quad (4.150)$$

for $0 < \beta < 1$. Inversion of the above Laplace transforms gives the functions,

$$\psi(t) = -\frac{d}{dt} \mathcal{E}_\beta(-t^\beta) \quad (4.151)$$

$$\Psi(t) = \mathcal{E}_\beta(-t^\beta) \quad (4.152)$$

where \mathcal{E}_β denotes the Mittag-Leffler function of order β . This transcendental function is defined through the power series,

$$\mathcal{E}_\beta(-t^\beta) = \sum_{n=0}^{\infty} \frac{(-1)^n t^{\beta n}}{\Gamma(\beta n + 1)}. \quad (4.153)$$

The Mittag-Leffler function generalises the exponential function ($\beta = 1$) and for $0 < \beta < 1$ it interpolates a stretched exponential and a power law (Raberto et.al.

(2002)),

$$\mathcal{E}_\beta(-t^\beta) \sim \begin{cases} e^{-\frac{t^\beta}{\Gamma(1+\beta)}} & t \rightarrow 0 \\ \frac{t^{-\beta}}{\Gamma(1-\beta)} & t \rightarrow \infty. \end{cases} \quad (4.154)$$

4.2.3 Empirical analysis of waiting times

In this section we will use high frequency ASX data to investigate the properties of the waiting time density and survival functions. As our data set consists of a tick-by-tick record of transactions we can directly observe the waiting time between consecutive trades. We will study the interval June 2000 to June 2001 for a number of different stocks, some traded with lower frequency than others, in order to build a picture of how the waiting time varies with liquidity. We will examine how well the assumptions of the CTRW hold up against real data.

The waiting time distribution is measured by differencing the time between consecutive trades. For each waiting time we also record the log-return for this transaction. The log-return becomes the jump value of the random walk. We show the waiting times for CBA during the studied period in Fig. 4.13 and we calculate the autocorrelation function for the waiting times. In the autocorrelation we observe a strong periodic trend present with period ~ 125 trades. This trend is reminiscent of the intraday volatility trends found in Chapter 3. Examining the histogram of the times at which trades took place we find an intraday pattern with most trades taking place in the morning or afternoon. The histogram shows the total number of transactions that occur during a given time interval, and is equivalent to the intraday pattern for the transaction frequency, $N_{t,\tau}$, in Chapter 3. We again assume that this behaviour is due to a deterministic time shift, and de-seasonalise the data by mapping to an appropriate business time scale. Fig. 4.13 displays the business time scale constructed such that the number of trades per unit time is, on average, a constant. Upon mapping the transaction times to this business time scale we find that the waiting times lose their periodic structure, as shown by the autocorrelation in Fig. 4.14.

A further assumption of the Montroll-Weiss CTRW is that the waiting times and the jump sizes are independent. By plotting the waiting time as a function of log-

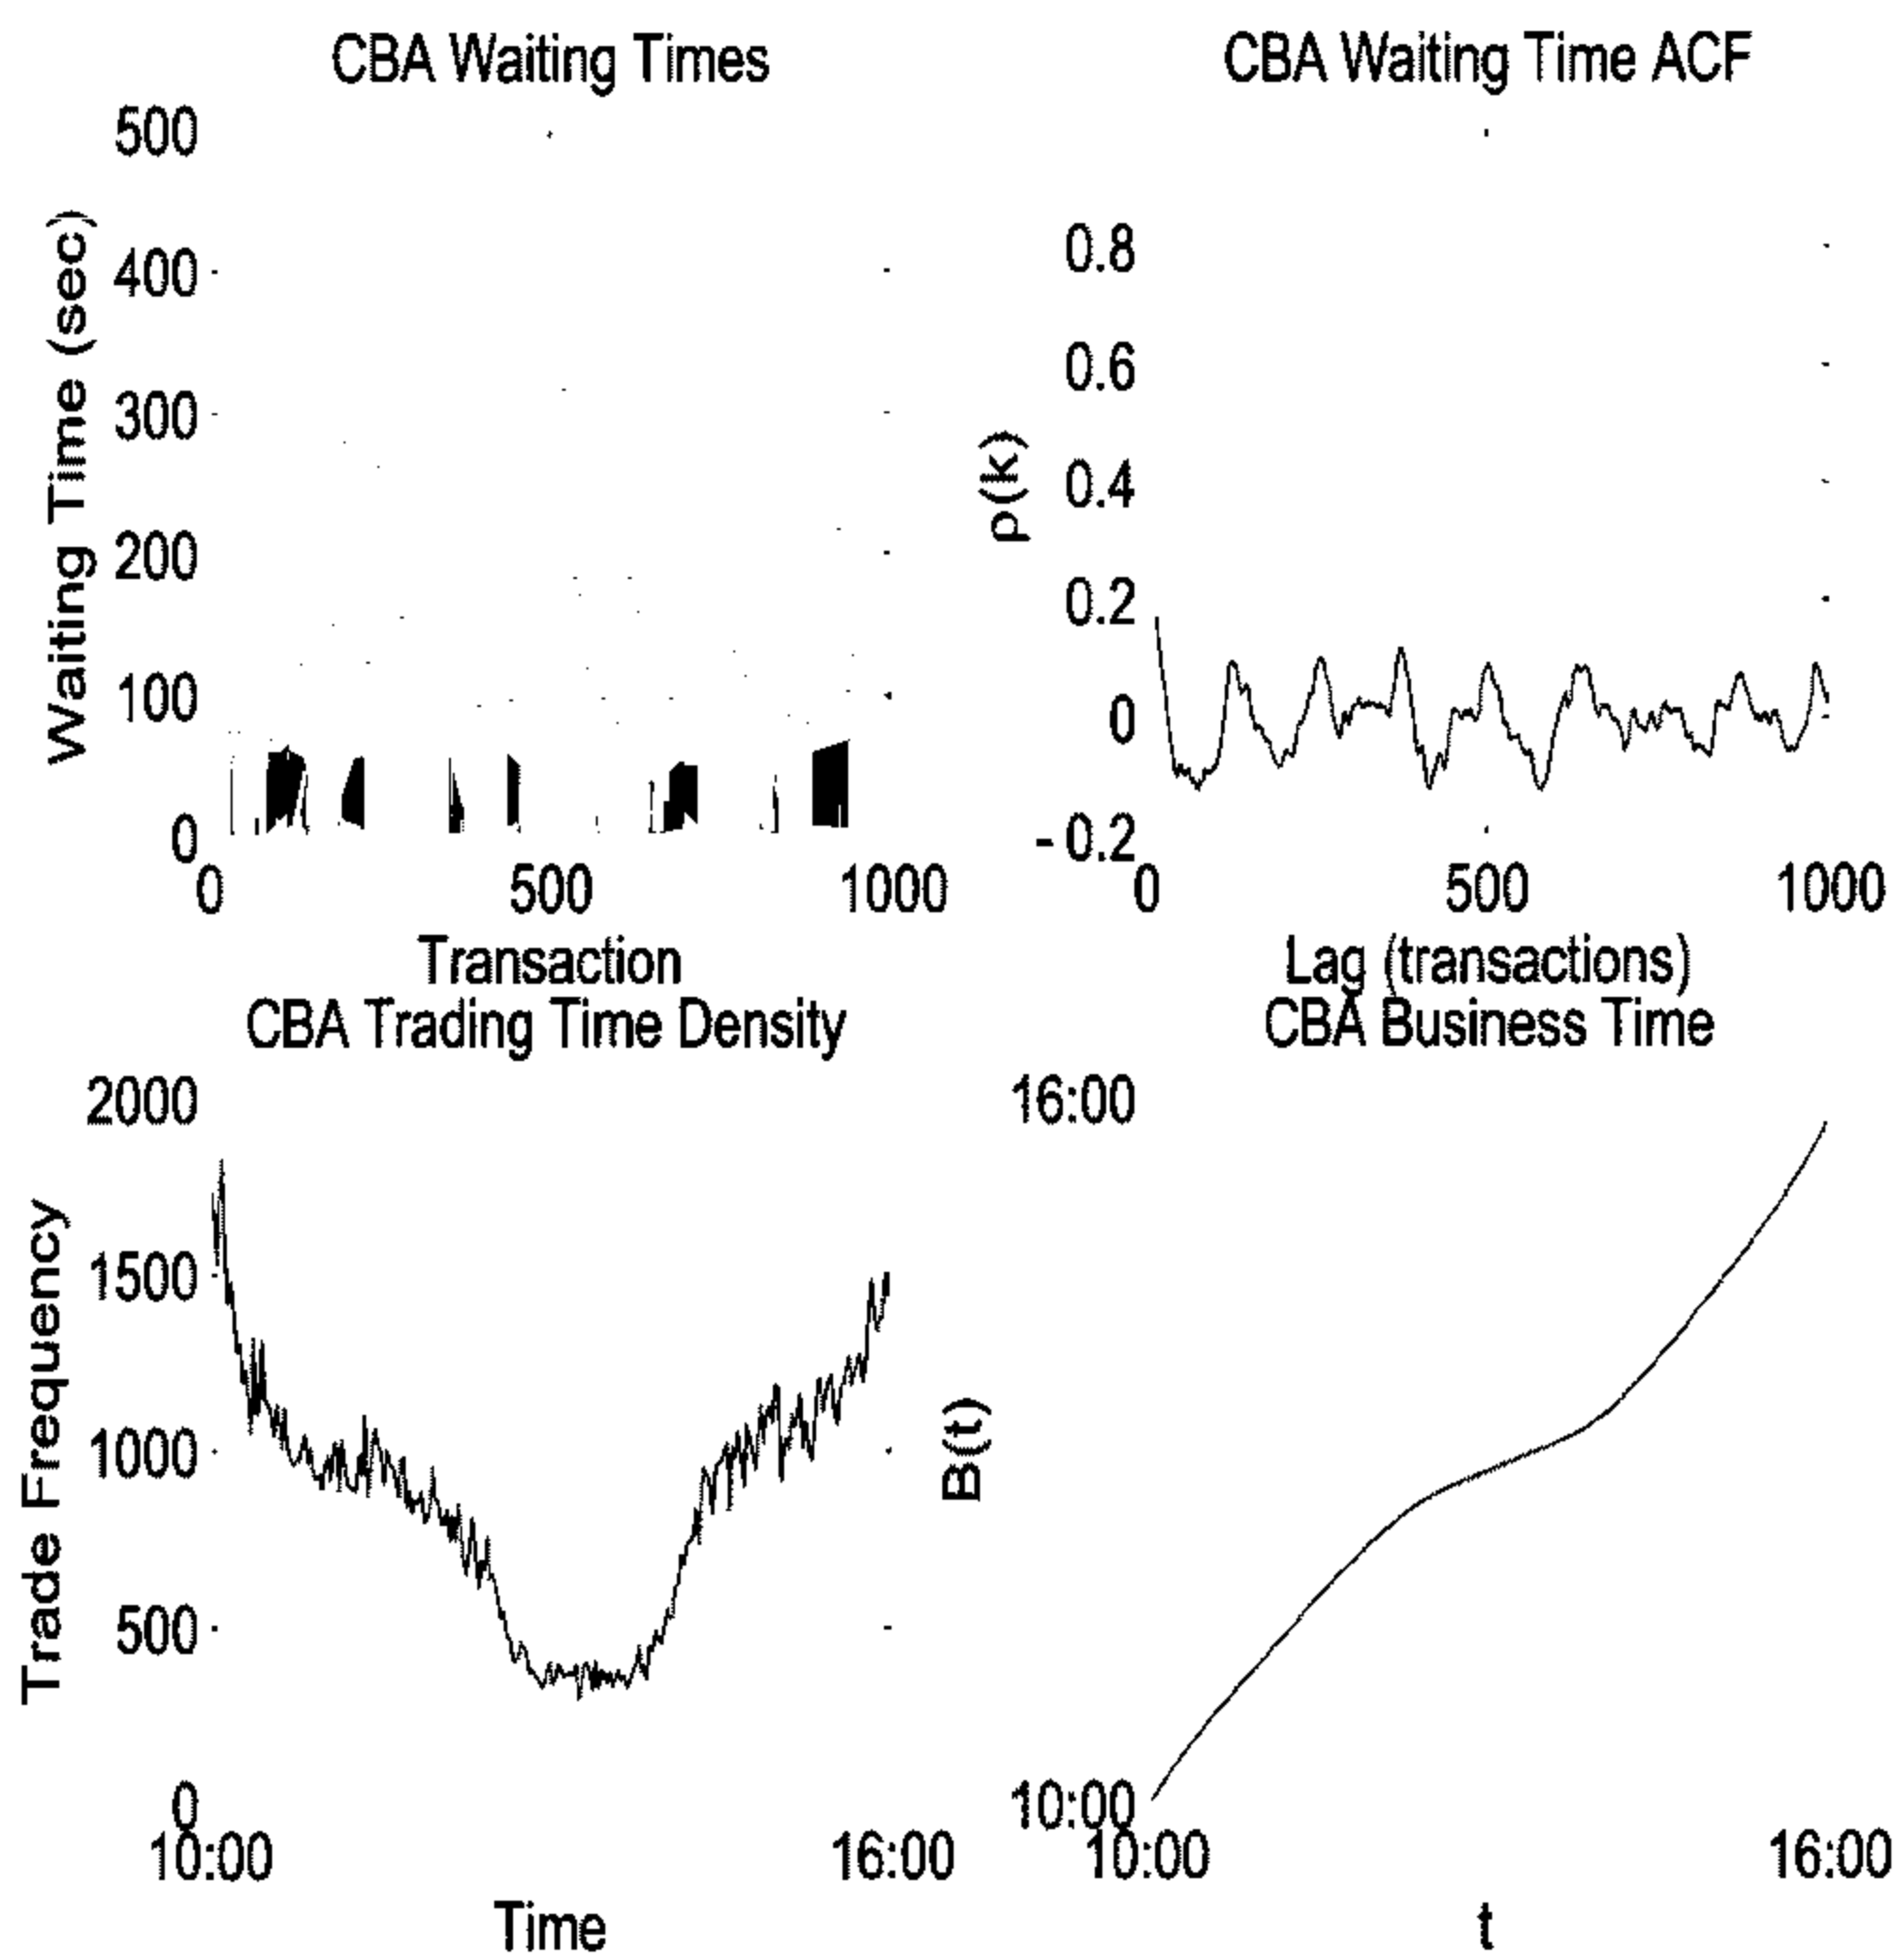


Figure 4.13: The top left plot displays the waiting times for CBA. The top right shows the ACF for the waiting times. The ACF shows a periodic structure. The plot in the lower left displays the histogram of trading times and indicates a periodic trend. The lower right displays the business time scale constructed from this trend.




Figure 4.14: The autocorrelation of the waiting times after mapping to business time. The periodic behaviour has successfully been removed.

return, in Fig. 4.15 we can see that there appears to be correlation between the two series. In Raberto et.al. (2002) a correlation was detected, however this may have been influenced by the periodic nature of their time series. By de-seasonalising our data we have the advantage that our results are not influenced by the deterministic intraday trend. Table C.2 presents the results of a hypothesis test for independence. We find that the hypothesis is rejected for most cases, except in the cases where τ_0 is relatively large. With this result we must consider the possibility of a more general Scher-Lax type uncoupled CTRW with evolution equation,

$$p(x, t) = \delta(x)\Psi(t) + \int_0^t \int_{-\infty}^{\infty} \phi(x - x', t - t')p(x', t')dx'dt'. \tag{4.155}$$

In Section 4.2.1 it was shown that if $\psi(t)$ is exponential then the CTRW is Markovian in nature. In fact if,

$$\psi(t) = \mu e^{-\mu t} \tag{4.156}$$

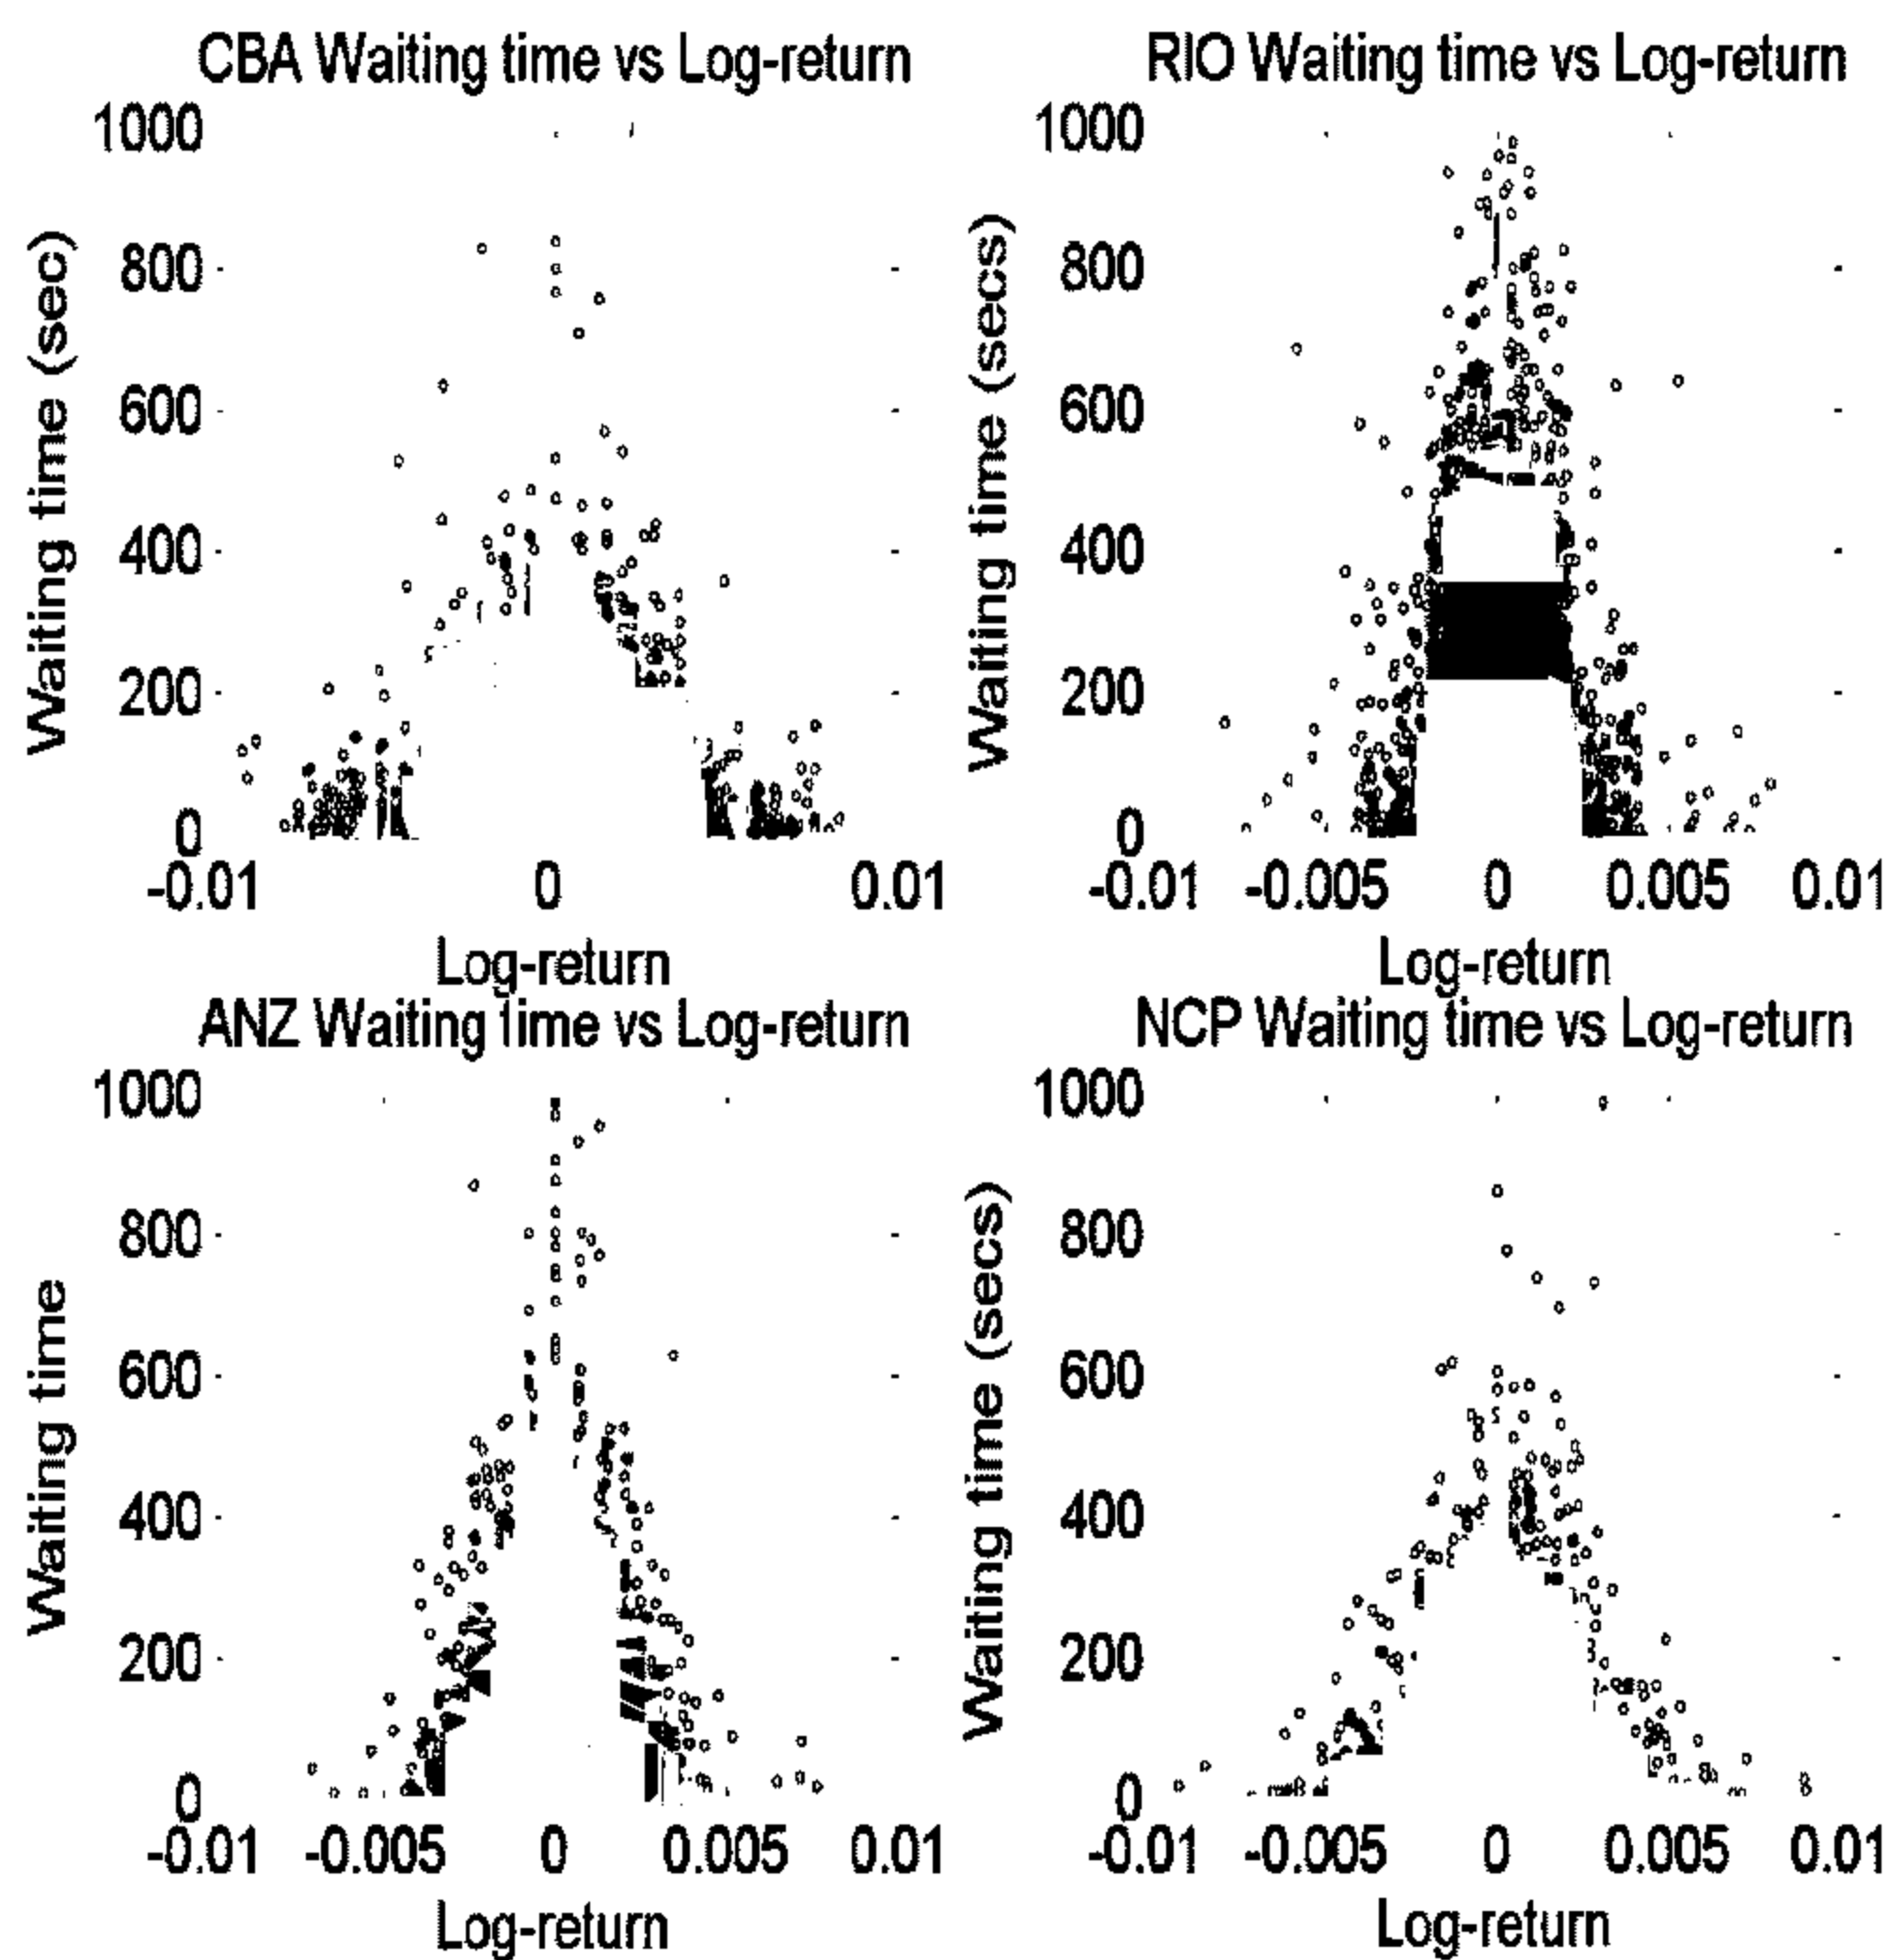


Figure 4.15: The waiting time plotted as a function of log-return for CBA, RIO, ANZ and NCP. If the two series were independent we would expect to see a uniformly scattered pattern. Instead we find a pattern that indicates short waiting times are more probable for large log-returns and large waiting times are more probable for small log-returns.

the probability that n jumps occur up to time t , $P(n, t)$ is given by the Poisson process,

$$P(n, t) = \frac{(\mu t)^n}{n!} e^{-\mu t}. \quad (4.157)$$

Hence the CTRW is a Poisson process with log-return density (Scalas et.al. (2004)),

$$p(x, t) = \sum_{n=0}^{\infty} \frac{(\mu t)^n}{n!} e^{-\mu t} \lambda_n(x). \quad (4.158)$$

With this choice of waiting time density, the CTRW is related to the compound Poisson models of Press (1976) and Rydberg and Shepard (1999). Many models of price formation assume that traders place their orders in an uncorrelated way. This assumption leads to order placement time being modelled by a Poisson process, as in Luckock (2003). If price formation were a thinning of the bid-ask spread then one should expect waiting times to be exponentially distributed. In Fig. 4.16 we display the complementary cumulative distribution, $\Psi(t)$ for the data. Under the assumption of market agent models such as in Luckock (2003) we would expect to find,

$$\Psi(\tau) \sim e^{-\frac{\tau}{\tau_0}} \quad (4.159)$$

where τ_0 is the empirical average of waiting times. As seen in Fig. 4.16 the theoretical exponential survival function does not offer an accurate representation for the empirical distribution. This result casts serious doubt on the assumptions of market agent models, and may indicate that time correlations are present at the bid-ask level. There are currently several studies taking place to investigate this possibility.

In Fig. 4.17 we show the result of fitting the stretched exponential function,

$$\Psi(\tau) \sim e^{-\left(\frac{\tau}{\tau_0}\right)^\beta} \quad (4.160)$$

to the empirical survival function. As the waiting time τ becomes larger, the distribution tends to flatten out according to a power law. We find that this function is a good fit to the empirical data for many different stocks in the data set. This behaviour corresponds to the memory function of Eq. (4.144) for small waiting times $\tau \rightarrow 0$. Table C.3 displays the goodness of fit, as measured by Pearson's r ,

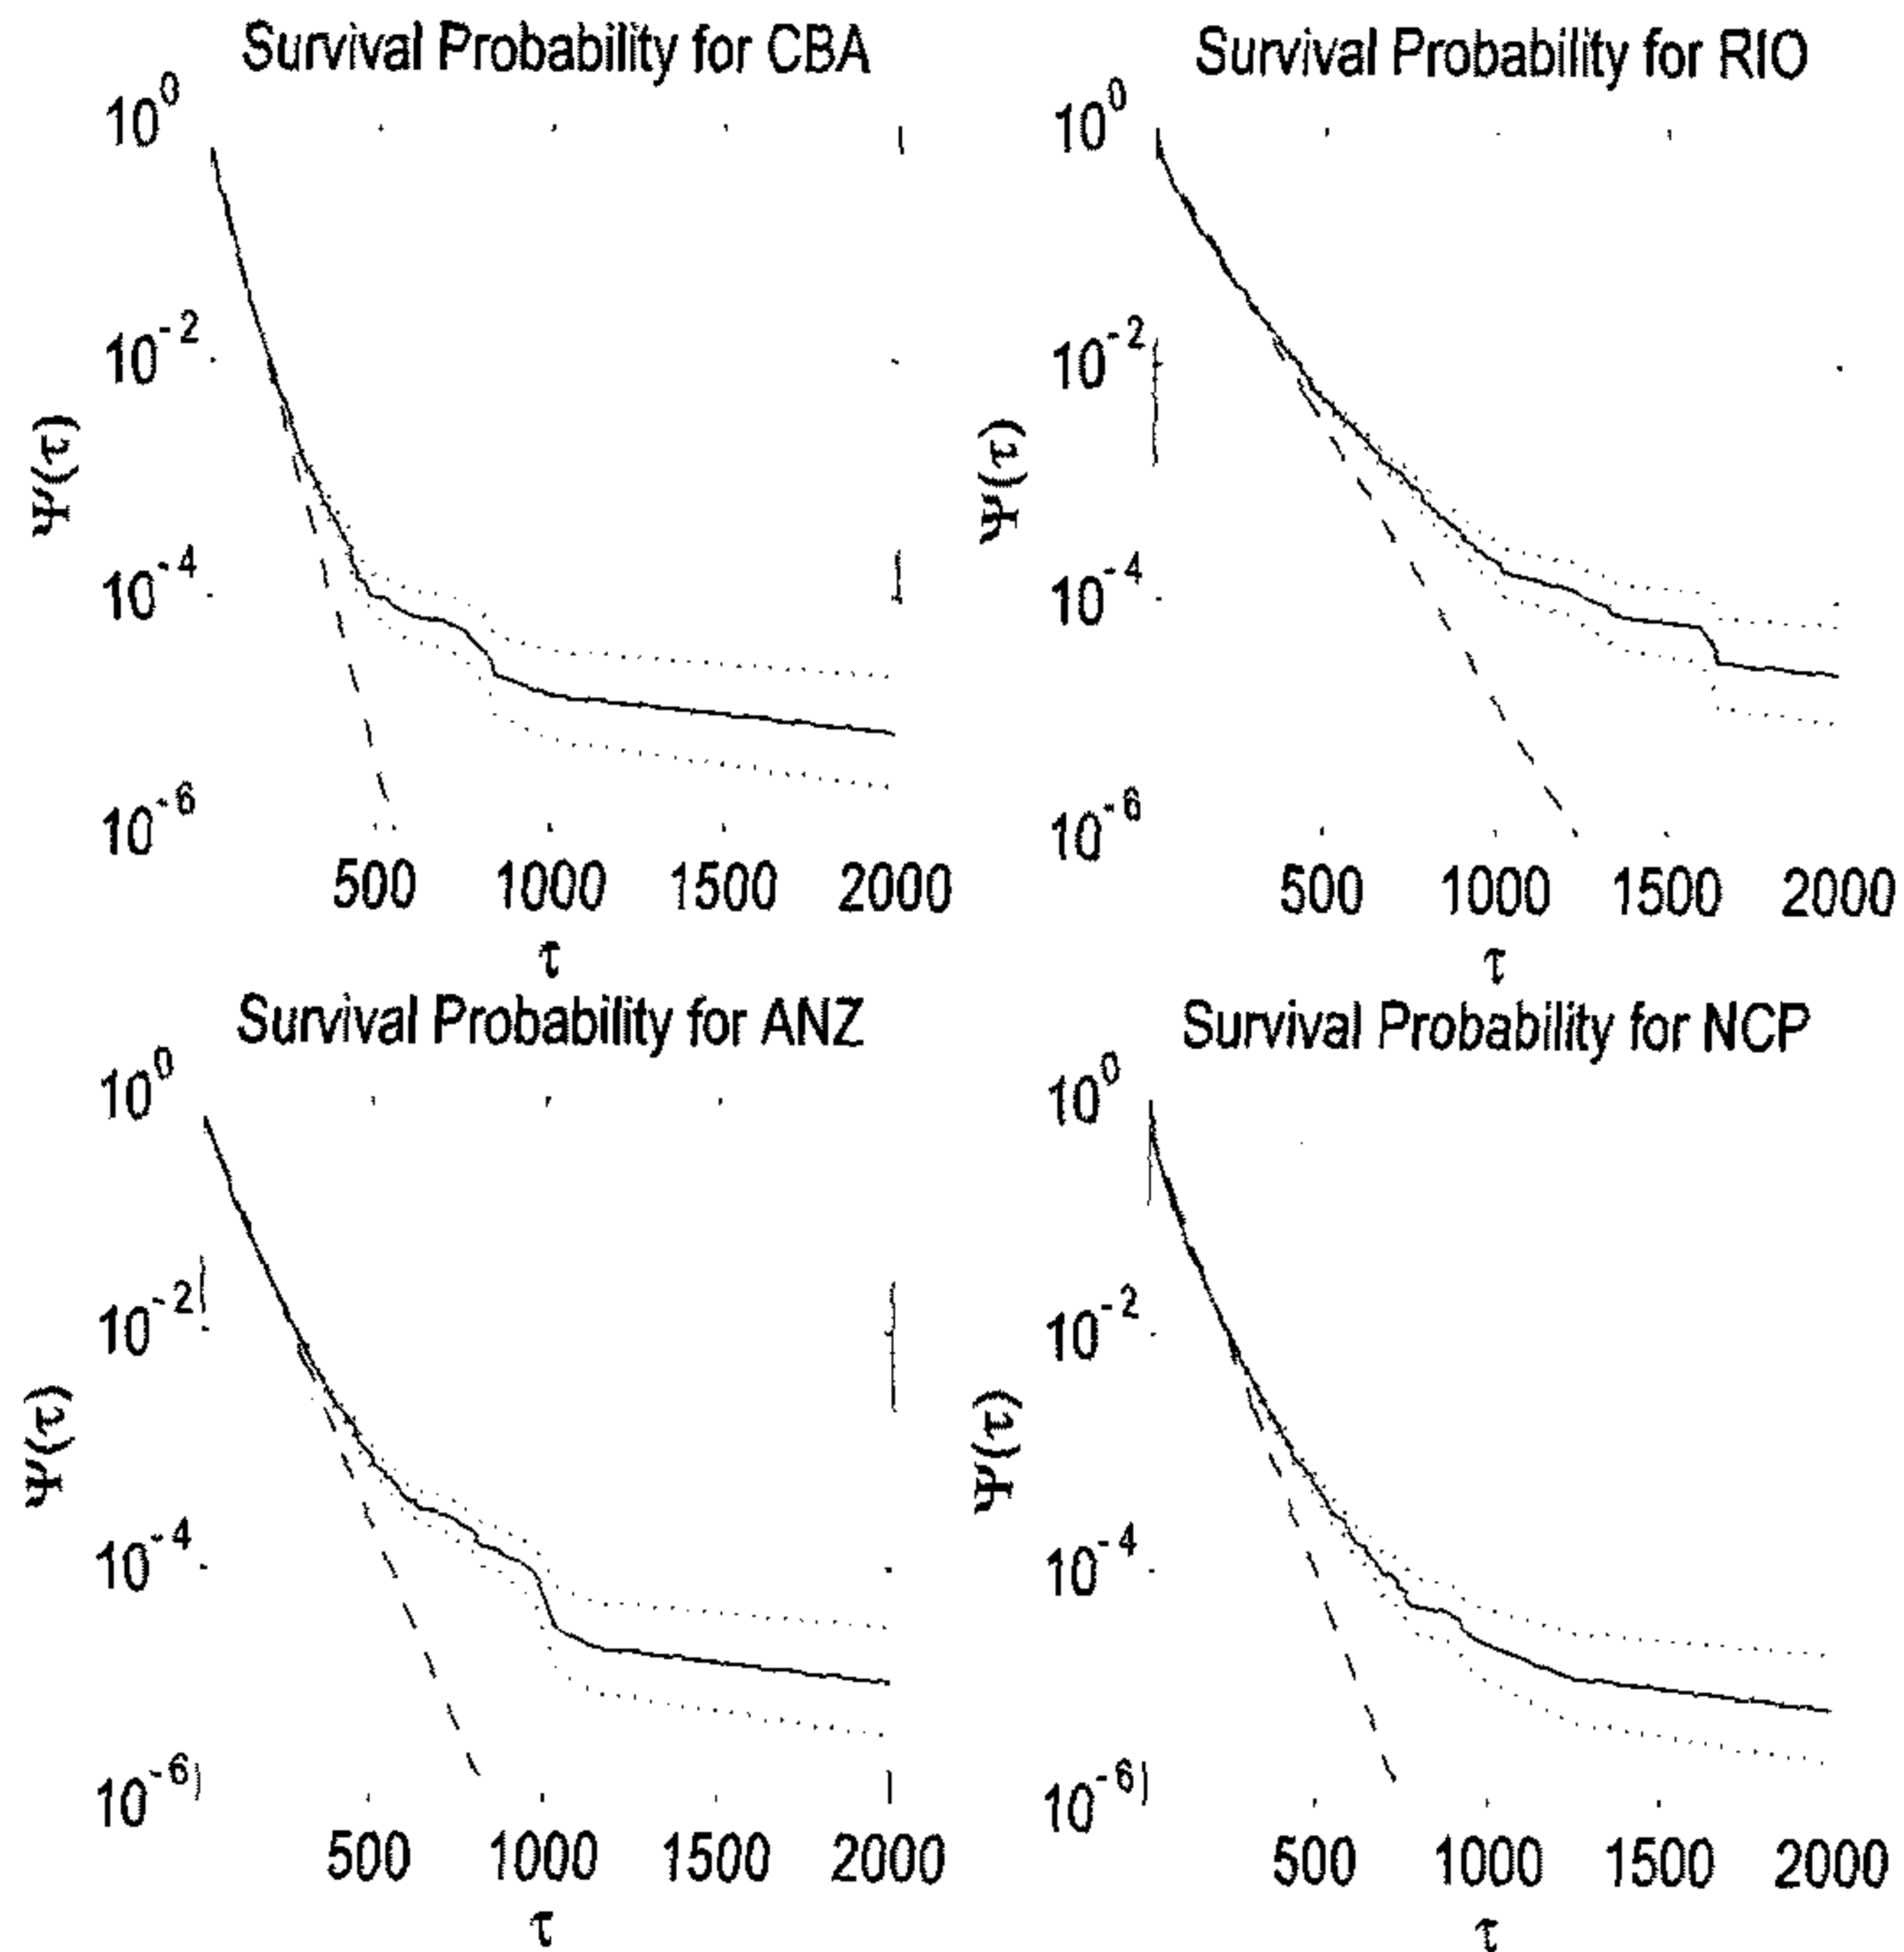


Figure 4.16: This plot displays the result of fitting the empirical survival function (**solid line**) with exponential distribution (**dashed line**). Agent based models are based on the assumption of exponential waiting times, however the poor fit displayed here casts serious doubt on the validity of this assumption.

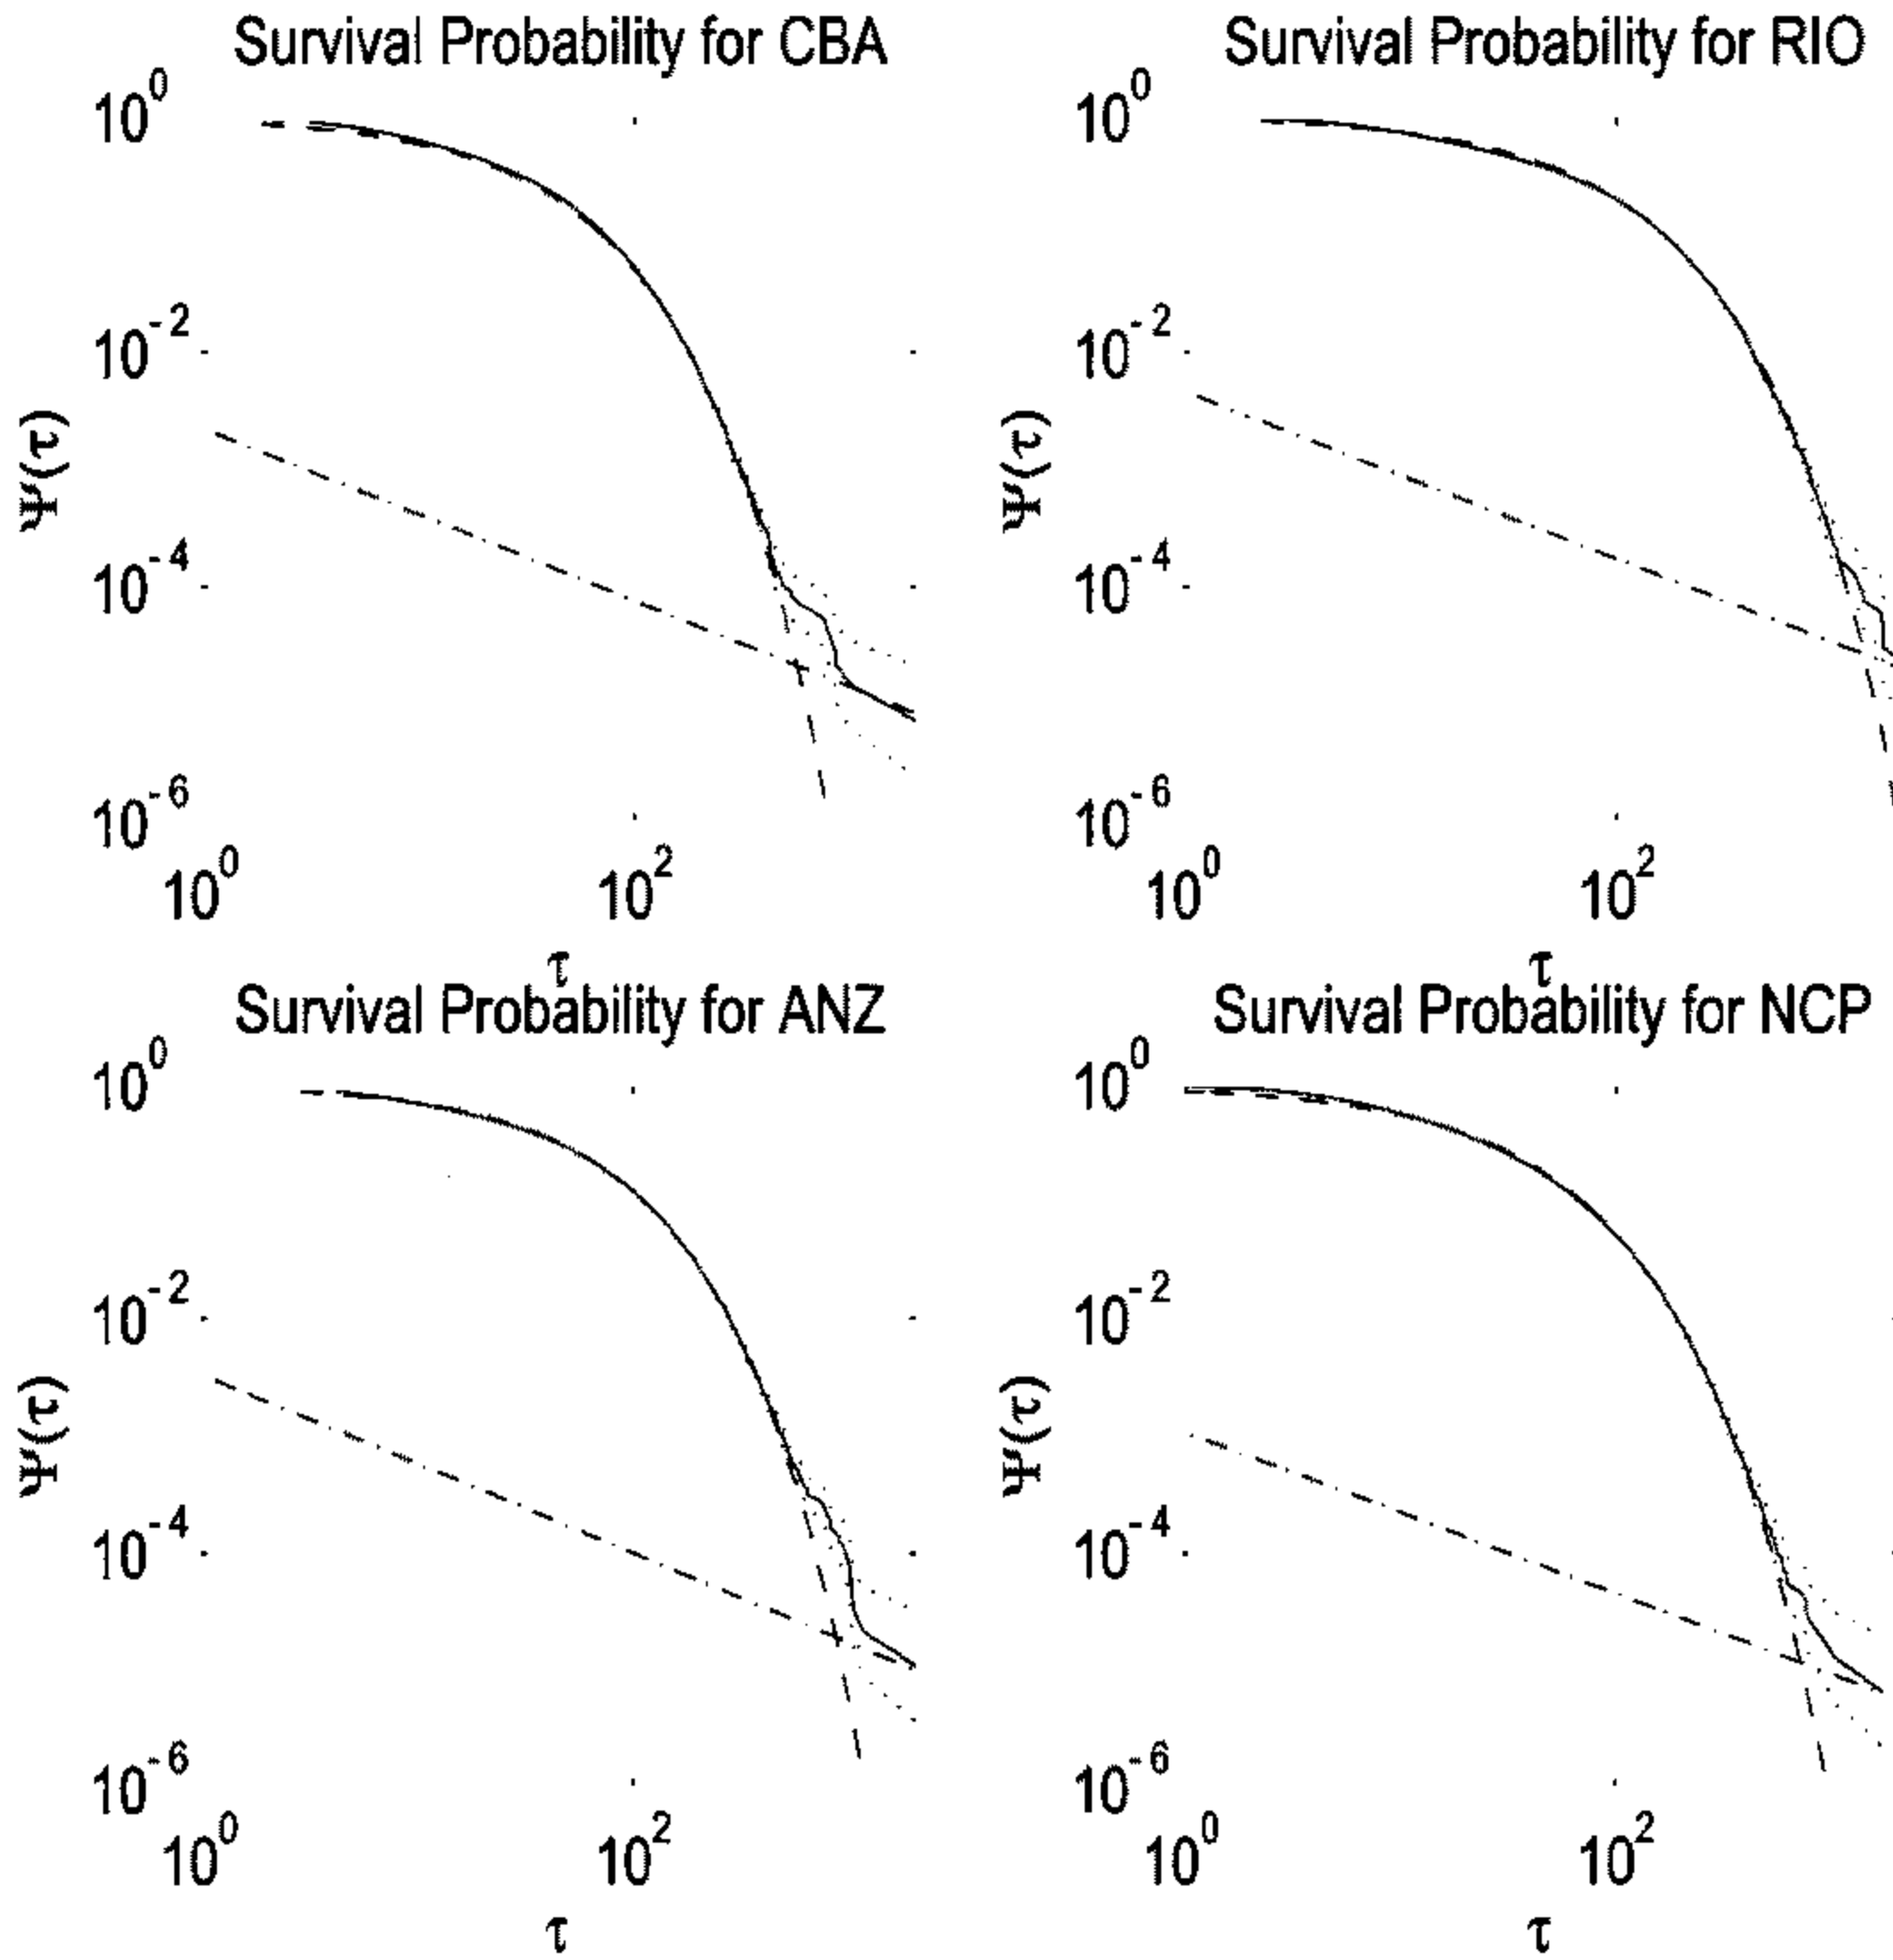


Figure 4.17: This plot displays the result of fitting the empirical survival function (solid line) with the Mittag-Leffler function of order β . The stretched exponential (dashed line) can be seen as offering a good approximation to the small τ part of the curve. The power-law behaviour (dot-dashed line) represents the tail behaviour of the survival function.

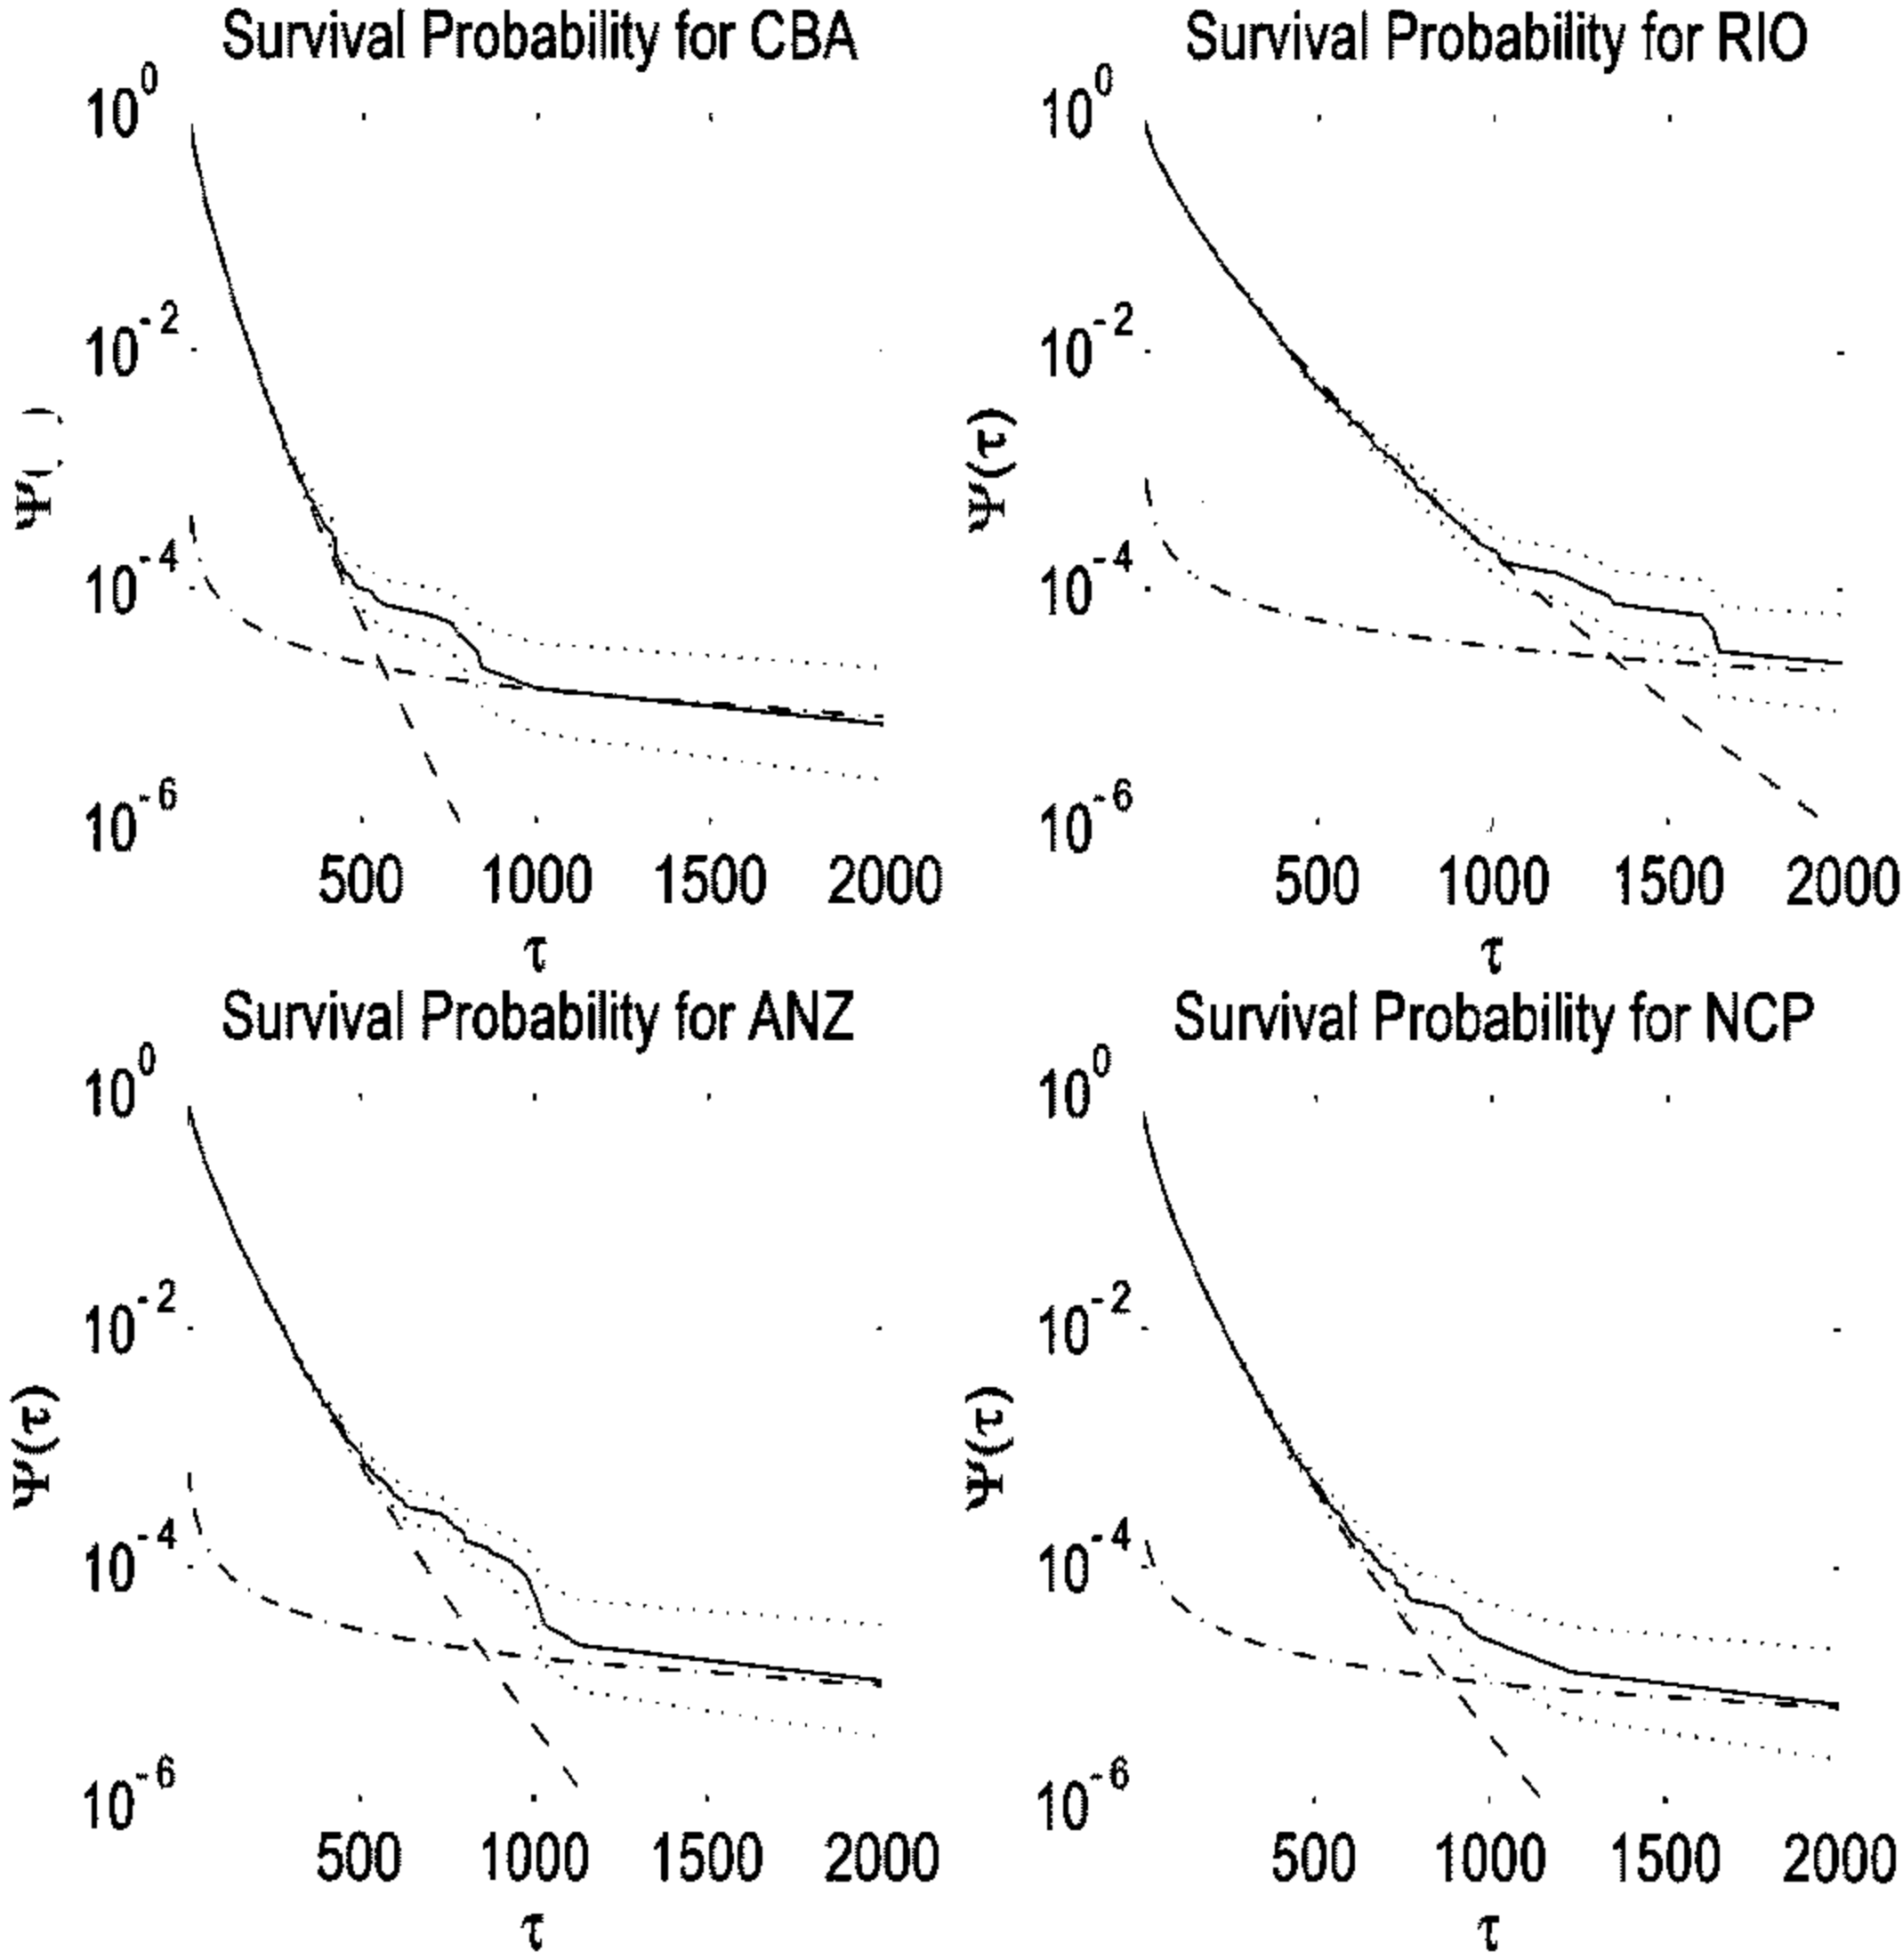


Figure 4.18: This plot displays the result of fitting the empirical survival function (solid line) with the Mittag-Leffler function of order β . The stretched exponential (dashed line) can be seen as offering a good approximation to the small τ part of the curve. The power-law behaviour (dot-dashed line) represents the tail behaviour of the survival function.

for the exponential and stretched exponential. We find that the stretched exponential consistently out performs the exponential function. We also plot the power law behaviour of Eq. (4.154) with the same value of β . The semi-log plot of Fig. 4.18 indicates that the survival function approaches this power law behaviour for large τ . Listed in Table C.3 are the average waiting times for each stock. In this study the waiting times τ are found to be smaller than 200 seconds, justifying the stretched exponential approximation to the Mittag-Leffler function. This finding supports the results of another study, Raberto et.al. (2002), who examined the waiting times for a selection of stocks trading on the New York Stock Exchange. Another study, Mainardi et.al. (2000), showed that for LIFFE BUND futures, waiting times could be in excess of 10,000 seconds and were well approximated by a Mittag-Leffler function with power law decay.

4.3 Summary

In this chapter we have examined some recent models for high frequency price dynamics. We proposed a new model for asset dynamics in a first attempt to describe the behaviour observed on the ASX. Our model was developed in order to explain the phenomenon of zero return enhancement. We termed the model threshold Brownian motion in reference to the market activity process, A_t , that must be above a threshold level, a , in order to produce a price change. The distribution properties of the log-return under TBM were shown to be heavily dependent on the properties of the occupation time of the process A_t . We looked at the properties of the model for several examples of activity process, two of which were used to represent the small τ and large τ behaviour of the market activity.

For small time intervals we approximated the behaviour of the activity with a Wiener process. The occupation time distribution for the Wiener process has a large probability for $T_{t,\tau} = 0$. This means that for this choice of activity process, TBM will specifically model zero return enhancement. The results of fitting the model to data show that the model is able to reproduce the large zero return spike found in the ASX data.

For large time intervals we do not see specific occurrences of zero return enhancement in ASX data. In this case by modelling the activity with an Ornstein-Uhlenbeck process we do not specifically model large scale zero return enhancement. The nature of the occupation time density for this process is such that there is a low probability for $T_{t,\tau} = 0$ and $T_{t,\tau} = \tau$. Thus while not specifically allowing for zero returns on a large scale, the effect of small scale zero return enhancement is still present in the model. The results of fitting this model to the 1 hour log-return show that even though zero return enhancement is occurring on a small scale, the large scale log-return distribution is significantly affected by this. The shape of the resulting distribution differs markedly from a Gaussian and is an excellent representation of the log-return distribution. This results indicates that we cannot simply ignore the effect of zero return enhancement when analysing financial data.

The model focused on the central region of the log-return distribution, ignoring the behaviour of the distribution of the tails. It is seen in the results of Fig. 4.11 and Fig. 4.12 that this model implicitly gives the log-return distribution fatter tails and that these tails are able to represent accurately those of the log-return. While example processes chosen to illustrate the models behaviour provided an excellent representation of the log-return distribution, we note that they do possess two disadvantages. Firstly, they can allow for negative activity. It is possible to avoid this by looking at the magnitude of the activity, *i.e.* with $\mathbf{1}(|A_t| > a)$, or by choosing a large mean reversion level/rate for the OU-process. Secondly, in fitting the model to data we have chosen to ignore the long memory properties of the financial time series, although Eqs. (4.12-4.14) show how such behaviour may be incorporated into the model.

The CTRW model of Scalas et.al. (2000) was examined as a possible choice for modelling high frequency ASX equity prices. We reviewed the theory of continuous time random walks and tested the model using high frequency stock price data. Analysis of the ASX data indicated that the assumption of independence between log-returns and waiting times could not be justified. We found that the waiting times for ASX equities were well represented by the stretched exponential function. This result sides in favour of the work of Scalas et.al. (2000), Mainardi et.al. (2000)

and Raberto et.al. (2002) suggesting the Mittag-Leffler function as an appropriate choice of survival function under the non-Markovian CTRW with power law memory kernel.

Chapter 5

Conclusion and closing remarks

In this thesis we presented a comprehensive study on the subject of modelling stock price dynamics. We examined data from the Australian Stock Exchange and investigated the ability of various models to describe the data. The ability to represent asset dynamics properly is of great importance to many areas in finance such as risk management, option pricing and portfolio analysis. This subject is continually attracting interest from researchers of many different backgrounds. Most recently, the advent of electronic trading and high frequency financial data has led to the creation of the field of econophysics: the science of empirical finance. By studying high frequency data we have been able to reveal important qualitative properties and thereby obtain a clearer understanding of the nature of financial markets. A complete knowledge of the statistical properties of asset return dynamics is crucial for the purpose of building and testing models for price change.

This thesis has contributed to the emerging field of econophysics by studying the behaviour and properties of stocks traded on the Australian Stock Exchange. The results from this study are important in that they show that ASX behaviour differs significantly from the reported behaviour of other markets. We examined the history of asset modelling starting with the early developments of Bachelier, the proposal of geometric Brownian motion, the contributions of Mandelbrot and many of the subordinated models that followed. Using ASX daily price data we tested the ability of these early models to describe the unconditional log-return distribution. We also

examined the performance of these models using high frequency data. We studied the distribution, correlation and scaling properties of high frequency ASX equity data. This study contributed to the knowledge of the fundamental process of price change. We examined new models for high frequency price dynamics. We showed that these new models were able to offer explanations as to the various phenomena observed in ASX data that previous models could not.

In Chapter 2 we undertook a review of the most important models for price dynamics and examined their various properties. We compared the models of GBM, Mandelbrot and Taylor (1963), Clark (1973), Praetz (1972) and Madan and Seneta (1990) using ASX daily price data. This comparison found that the Student- t model of Praetz (1972), with $\nu \approx 4$, was the most successful model for describing the unconditional daily return. We examined several more advanced models that incorporated a correlation structure in their descriptive properties. Among these models was the FATGBM of Heyde (1999). By constructing the activity time we were able to investigate the FATGBM assumption of $T_t - t$ self-similarity with Australian daily data. The discrete time GARCH model was reviewed and its ability to describe ASX daily data was examined. We investigated a misclassification problem in discrete time series modelling whereby an ARMA model is unable to distinguish between different time series. The GAR(1) model of Peiris (2004) was shown to offer a possible solution to this problem and thereby increase forecast performance.

In Chapter 3 we looked at the real world behaviour of stock prices through the use of high frequency financial data. It was shown how the microstructure of the Australian market influences the behaviour of the data. We proposed that the ASX returns can be considered as two separate processes representing the overnight return and the intraday return. A strong periodic trend was found to exist in the intraday absolute return, the volume and in the number of transactions. By defining a suitable new time scale, business time, on which to operate we were able to effectively remove this periodic behaviour. Using the method of principal component analysis we constructed a new measure for the intraday market volatility called the master volatility measure. Examination of the log-return distribution revealed the presence of zero return enhancement, that is, a disproportionate number of zero

returns in the time series. We found that these zero returns were mainly generated by the stock trading at a constant price. The level of zero return enhancement was also found to vary over the course of a trading day. By constructing a new time scale based on the master volatility measure we were able to make the level of ZRE constant across a trading hours. The unconditional models of Chapter 2 were compared using high frequency data. This comparison indicated that the variance gamma model of Madan and Seneta (1990) was best able to describe high frequency returns. We also looked at the ability of GARCH to provide a conditional model for ASX high frequency stock returns. It was shown that the exponentially decaying correlation structure of GARCH was insufficient to describe the data. In all cases it was found that the models could not reproduce the sharp central spike in the data caused by zero return enhancement. Empirical examinations found evidence of power law behaviour in the tails of the various distributions studied with index values, on average, $\alpha_L \approx 3.6$, $\alpha_{|L_{t,\tau}|} \approx 3.6$, $\alpha_{V_{t,\tau}} \approx 3.4$ and $\alpha_{N_{t,\tau}} \approx 3.0$. This finding also provided evidence that ASX stock returns possess finite variance. The power law behaviour was also found to be consistent for different values of the sampling interval τ . The behaviour we detected on the ASX differs substantially to that found in other markets. This is an important finding as it suggests that price behaviour can be market specific, and that models and trading systems developed for one market may not necessarily be appropriate for another. In this chapter we also found evidence of long range dependence in ASX equity data. Using several different methods, we found that while the log-return appeared uncorrelated, there was LRD present in the absolute log-return, volume, and number of transactions. The estimated values for the long memory parameter were, on average, $\beta_L \approx 1$, $\beta_{|L|} \approx 0.6$, $\beta_V \approx 0.3$ and $\beta_N \approx 0.25$. We investigated the data for scaling behaviour using several different techniques. By re-scaling empirical density functions it was shown that the data exhibits scaling behaviour. We used a wavelet based method to test for self-similarity and to estimate the scaling parameter, H , of the data. We explored further the nature of this behaviour by calculating Renyi exponents to test for multifractal behaviour. The Renyi exponents indicated that for high frequency Australian data, a linear scaling relationship was sufficient. This finding explains

the time invariance of the power law estimates of tail behaviour (Fig. 3.25).

In Chapter 4 we presented a new model for the behaviour of price dynamics. This model is able to explain the phenomenon of zero return enhancement (ZRE) through the employment of an underlying activity process, A_t , that represents the flow of market information. We showed that the behaviour of the model was dependent upon the occupation time of A_t . By calculating the occupation time for several different choices for A_t we were able to investigate the properties of the model. We fitted the model to high frequency data for two time scales. Over a small time scale, $\tau = 10$ minutes, we modelled the market activity as a Wiener process. The nature of the Wiener process' occupation time meant that the stock was most likely to trade all of the time or not trade at all. This behaviour provided specific instances of zero return enhancement and reproduced the central spike of the empirical density function. On a large time scale, $\tau = 1$ hour, we modelled the activity as an Ornstein-Uhlenbeck process. In this case, although the model did not show specific instances of ZRE, the periods of constant trading had the effect of fattening the log-return distribution. The model was found to offer good fits to the data for both cases. These results highlight the importance of the zero returns in the financial data. Zero return enhancement cannot simply be ignored as it has consequences for the shape of the distribution over large sampling intervals. Finally we studied the continuous time random walk (CTRW) model for high frequency price dynamics. We investigated the theory of the CTRW and showed the general evolution equations for the non-Markovian case. In the case where the memory function was specified by a function with power-law decay the survival probability and waiting time distribution are given in terms of the Mittag-Leffler distribution. We undertook an analysis of price data and waiting times in order to investigate the assumptions and behaviour of the CTRW model. The waiting times were found to display a strong periodic trend. This periodicity was successfully removed by mapping onto an appropriate new time scale similar to that in Chapter 3. We found that the waiting times and the log-return were not independent as is required by the Montroll-Weiss CTRW. We also found that the empirical survival function was well approximated by the Mittag-Leffler function. This finding casts serious doubt on the assumptions of many

models of price formation, where the waiting time distribution is often chosen to be exponentially distributed.

Despite more than 50 years of research, the fundamental nature of the process of price change still alludes researchers. This thesis has presented a study of the Australian market with the aim of contributing to the general understanding of financial markets. It is hoped that eventually the knowledge gained from studies such as this will lead to a more complete characterisation of the stochastic processes of price change and will lead to the ability to construct better models for price dynamics.

Appendix A

Tables for Chapter 2

Lévy parameters:	α	β	σ	μ
Security				
ANZ	1.73	-0.03	0.63	0.06
CBA	1.75	0.37	0.62	0.12
RIO	1.60	0.23	0.61	0.07
WMC	1.78	0.34	0.62	0.03

Table A.1: This table shows the estimated parameters for the unconditional four parameter Lévy-stable model. The parameters for the stocks shown here are found to be much the same.

Lévy:	α	σ	μ
Security			
AMC	1.5743	0.5579	0.0029
ANZ	1.6767	0.6273	0.0147
CBA	1.7457	0.6312	-0.0063
CML	1.5399	0.5281	0.0173
CSR	1.6196	0.5892	-0.0029
FGL	1.5750	0.4998	-0.0164
NAB	1.6998	0.6008	0.0317
NCP	1.5257	0.5367	-0.0282
RIO	1.6151	0.6171	-0.0463
SGB	1.6078	0.5740	-0.0113
SUN	1.2292	0.4916	0.0001
WBC	1.6070	0.6170	0.0295
WES	1.4587	0.5656	-0.0116
WFT	1.7246	0.6068	0.0118
WMC	1.7955	0.6209	-0.0422
WOW	1.5792	0.5958	0.0142
WPL	1.6322	0.5915	-0.0519

Table A.2: This table shows the estimated parameters for Mandelbrot and Taylor three parameter, subordinated Lévy-stable model. The parameters are found to be roughly consistent across all the stocks, with $\alpha \sim 1.7$, $\sigma \sim 0.6$ and $\mu \sim 0$.

GBM:	μ	σ
Security		
AMC	0.0060	0.7923
ANZ	0.0169	0.8899
CBA	-0.0015	0.8937
CML	0.0223	0.7485
CSR	-0.0014	0.8347
FGL	-0.0198	0.7117
NAB	0.0307	0.8507
NCP	-0.0253	0.7632
RIO	-0.0416	0.8771
SGB	-0.0041	0.8145
SUN	0.0022	0.7158
WBC	0.0226	0.8747
WES	-0.0102	0.8062
WFT	0.0123	0.8572
WMC	-0.0404	0.8771
WOW	0.0093	0.8488
WPL	-0.0463	0.8383

Table A.3: This table shows the estimated parameters geometric Brownian motion. The parameters indicate that for all the stocks, $\mu \sim 0$ and $\sigma \sim 0.8$.

Clark:	μ	σ	χ
Security			
AMC	3.4985	0.2741	1.2823
ANZ	3.4709	0.2951	1.1604
CBA	3.5686	0.2799	1.0536
CML	3.6510	0.2470	1.2326
CSR	3.5446	0.2752	1.1844
FGL	3.3562	0.2983	1.3157
NAB	3.6371	0.2617	1.0443
NCP	3.6084	0.2546	1.2646
RIO	3.3212	0.3136	1.2721
SGB	3.5181	0.2755	1.2340
SUN	3.2924	0.3763	1.7543
WBC	3.6135	0.2701	1.0523
WES	3.5475	0.3809	1.6000
WFT	3.6239	0.2652	1.0567
WMC	3.6775	0.2517	0.8677
WOW	3.0282	0.3402	1.4356
WPL	3.5080	0.2814	1.2062

Table A.4: This table shows the estimated parameters for Clark's subordinated lognormal model. Like the Mandelbrot and Taylor (1967) and GBM models, the parameters for Clark's model seem to be consistent across stocks. The nature of the driving process means that the parameter μ is much larger than for the other models.

VG:	μ	σ	v
Security			
AMC	0.0074	0.9159	0.4526
ANZ	0.0114	0.9965	0.3625
CBA	-0.0126	0.9804	0.2956
CML	0.0174	0.8524	0.3939
CSR	-0.0028	0.9441	0.3833
FGL	-0.0025	1.0023	1.0372
NAB	0.0323	0.9401	0.3140
NCP	-0.0249	0.8843	0.4592
RIO	-0.0459	1.0055	0.4378
SGB	-0.0187	0.9326	0.4265
SUN	0.0026	0.9877	0.9852
WBC	0.0341	0.9822	0.3575
WES	-0.0006	1.0955	0.9383
WFT	0.0016	1.1906	1.1433
WMC	-0.0426	0.9320	0.1877
WOW	0.0208	1.0294	0.6316
WPL	-0.0520	0.9547	0.4085

Table A.5: This table shows the estimated parameters for the variance gamma model of Madan and Seneta (1990). For this model, $\mu \sim 0$, $\sigma \sim 1$ but v appears to vary between stocks.

Student- t :	μ	σ	ν
Security			
AMC	0.0038	0.7211	3.9125
ANZ	0.0132	0.8281	5.2683
CBA	-0.0092	0.8461	6.8639
CML	0.0168	0.6794	3.6750
CSR	-0.0029	0.7683	4.3859
FGL	-0.0151	0.6406	3.6973
NAB	0.0319	0.8002	6.0048
NCP	-0.0270	0.6850	3.4339
RIO	-0.0466	0.8015	4.2683
SGB	-0.0139	0.7473	4.3215
SUN	0.0004	0.5687	1.7573
WBC	0.0319	0.8059	4.4301
WES	-0.0113	0.7007	2.7124
WFT	0.0117	0.8134	6.6951
WMC	-0.0428	0.8456	9.6959
WOW	0.0173	0.7630	3.6886
WPL	-0.0529	0.7742	4.6291

Table A.6: This table shows the estimated parameters for the Student- t model of Praetz (1972). The Student- t model has finite moments of order k , where $k < \nu$. The estimated parameters indicate that for most stocks, $\nu > 4$. This indicates a finite mean, variance and kurtosis.

Models:	GBM	Lévy	Clark	VG	Student- <i>t</i>
Security					
AMC	-3.32×10^3	-3.00×10^3	-3.102×10^3	-2.99×10^3	-2.98×10^3
ANZ	-3.22×10^3	-3.17×10^3	-3.16×10^3	-3.15×10^3	-3.14×10^3
CBA	-3.21×10^3	-3.14×10^3	-3.16×10^3	-3.14×10^3	-3.13×10^3
CML	-3.41×10^3	-2.94×10^3	-3.14×10^3	-2.97×10^3	-2.91×10^3
CSR	-3.27×10^3	-3.11×10^3	-3.16×10^3	-3.10×10^3	-3.08×10^3
NAB	-3.25×10^3	-3.09×10^3	-3.17×10^3	-3.10×10^3	-3.08×10^3
NCP	-3.39×10^3	-2.97×10^3	-3.15×10^3	-2.98×10^3	-2.95×10^3
RIO	-3.23×10^3	-3.17×10^3	-3.14×10^3	-3.14×10^3	-3.14×10^3
SGB	-3.30×10^3	-3.06×10^3	-3.13×10^3	-3.06×10^3	-3.04×10^3
SUN	-3.49×10^3	-3.01×10^3	-2.96×10^3	-2.93×10^3	-2.97×10^3
WBC	-3.23×10^3	-3.19×10^3	-3.21×10^3	-3.15×10^3	-3.15×10^3
WES	-3.27×10^3	-3.10×10^3	-3.05×10^3	-3.06×10^3	-3.07×10^3
WFT	-3.24×10^3	-3.12×10^3	-3.18×10^3	-3.13×10^3	-3.11×10^3
WMC	-3.23×10^3	-3.11×10^3	-3.189×10^3	-3.13×10^3	-3.11×10^3
WOW	-3.26×10^3	-3.13×10^3	-3.10×10^3	-3.11×10^3	-3.11×10^3
WPL	-3.27×10^3	-3.11×10^3	-3.16×10^3	-3.10×10^3	-3.09×10^3

Table A.7: This table shows the log-likelihood values for the models estimated in Tables A.2-A.6. The Student-*t* model of Praetz displays the largest (least negative) values, followed closely by the variance gamma model. All the models outperform GBM.

Parameters:	Volume			Transaction frequency		
	GAR(1)		AR(1)	GAR(1)		AR(1)
	α_{GAR}	δ	α_{AR}	α_{GAR}	δ	α_{AR}
BIL	0.85	0.80	0.85	0.75	0.74	0.85
BHP	0.83	0.70	0.81	0.87	0.73	0.79
CBA	0.70	0.64	0.75	0.78	0.78	0.76
CML	0.42	0.89	0.40	0.74	0.91	0.70
RIO	0.54	0.46	0.47	0.83	0.52	0.83
WMC	0.50	0.81	0.47	0.67	0.99	0.70

Table A.8: This table displays the parameters obtained for the GAR(1) and AR(1) models using volume and transaction data from the Australian Stock Exchange.

Forecasts	Volume		
	GAR(1)	AR(1)	Observed
$t = 1001$	0.49	0.65	0.51
$t = 1002$	0.30	0.45	0.36
$t = 1003$	0.21	0.31	0.28

Table A.9: Shown are the forecasts for $t = 1001, 1002, 1003$ obtained with both AR(1) and GAR(1) models.

Appendix B

Tables for Chapter 3

Variables	Component number		
	1	2	3
$ L_t $	0.5738	0.6692	0.4722
V_t	0.5706	-0.7402	0.3556
N_t	0.5875	0.0654	-0.8066
Percent of total variation explained:	95.83%	3.80%	0.36%

Table B.1: Principal components based on the logarithm of the sample volatility, volume, and the transaction frequency.

	Scaling		
	$\mathbb{E}[L_t]$	$\mathbb{E}[V_t]$	$\mathbb{E}[N_t]$
Security			
AGL	0.5	1	1
ANZ	0.5	1	1
CBA	0.5	1	1
RIO	0.5	1	1
NCP	0.5	1	1
WBC	0.5	1	1
FGL	0.5	1	1
CSR	0.5	1	1
BIL	0.5	1	1
SGB	0.5	1	1

Table B.2: Values for scaling of for $\mathbb{E}[|L_t|]$, $\mathbb{E}[V_t]$ and $\mathbb{E}[N_t]$.

Lévy:	α	σ	μ
Security			
AGL	1.0924	0.5357	-0.0259
ANZ	1.4871	0.6503	0.0152
CBA	1.3718	0.6723	-0.0137
RIO	1.4232	0.4994	0.0095
NCP	1.3159	0.4472	0.0125
WBC	1.4041	0.6211	-0.0066
FGL	1.7955	0.7329	-0.0032
CSR	1.6129	0.657	-0.0009
BIL	1.158	0.5012	-0.028
SGB	1.0273	0.5428	-0.0111

Table B.3: This table shows the estimated parameters for Mandelbrot and Taylor three parameter, subordinated Lévy-stable model using high frequency data.

TLF:	σ	α	ζ
Security			
AGL	0.9784	1.0129	0.2719
ANZ	0.8683	1.3468	0.4842
CBA	0.8338	1.2993	0.3714
RIO	0.4018	1.3912	0.0568
NCP	0.6806	1.2196	0.392
WBC	0.9775	1.1022	0.463
FGL	1.0325	1.3598	0.5834
CSR	1.0087	1.0237	0.5137
BIL	0.9076	1.074	0.3778
SGB	0.9802	1.0115	0.2647

Table B.4: This table shows the estimated parameters for the truncated Lévy flight using high frequency data.

GBM:	μ	σ
Security		
ANZ	0.0196	0.8728
CBA	-0.0135	0.9129
RIO	0.011	0.6339
NCP	0.0156	0.5879
WBC	-0.0095	0.8003
FGL	-0.0044	1.0255
CSR	-0.0108	0.8918
BIL	-0.0442	0.596
SGB	-0.0116	0.558

Table B.5: This table shows the estimated parameters geometric Brownian motion using high frequency data.

Clark:	μ	σ	χ
Security			
AGL	3.4985	0.2741	1.2823
ANZ	3.4709	0.2951	1.1604
CBA	3.5686	0.2799	1.0536
BIL	3.6510	0.2470	1.2326
CSR	3.5446	0.2752	1.1844
FGL	3.3562	0.2983	1.3157
NCP	3.6084	0.2546	1.2646
RIO	3.3212	0.3136	1.2721
SGB	3.5181	0.2755	1.2340
WBC	3.6135	0.2701	1.0523

Table B.6: This table shows the estimated parameters for Clark's subordinated lognormal model using high frequency data.

VG:	μ	σ	ν
Security			
AGL	-0.0188	1.1266	2.0147
ANZ	0.0145	1.0249	0.615
CBA	-0.014	1.0784	0.8292
RIO	0.0083	0.8393	0.884
NCP	0.0121	0.8223	1.4389
WBC	-0.0045	1.0218	0.7869
FGL	-0.0035	1.0781	0.1869
CSR	0.0006	1.0046	0.3927
BIL	-0.0228	0.9739	1.6712
SGB	-0.0081	1.1444	2.0642

Table B.7: This table shows the estimated parameters for the variance gamma model of Madan and Seneta (1990) using high frequency data.

Student- t :	μ	σ	ν
Security			
AGL	-0.0263	0.5584	1.2726
ANZ	0.0162	0.7886	3.1993
CBA	-0.0138	0.7763	2.3586
RIO	0.0094	0.5926	2.5354
NCP	0.0116	0.5056	1.9482
WBC	-0.0067	0.7335	2.5984
FGL	-0.0034	0.9822	9.2182
CSR	-0.0007	0.8358	4.6161
BIL	-0.0281	0.5364	1.4615
SGB	-0.011	0.5486	1.1022

Table B.8: This table shows the estimated parameters for the Student- t model of Praetz (1970) using high frequency data.

Models:	GBM	Lévy	TLF	Clark	VG	Student- t
Security						
AGL	-3.1797×10^4	-2.3324×10^4	-2.1017×10^4	-2.1807×10^4	-2.0996×10^4	-2.3015×10^4
ANZ	-2.5911×10^4	-2.5125×10^4	-2.4529×10^4	-2.4542×10^4	-2.4439×10^4	-2.4737×10^4
CBA	-2.5701×10^4	-2.5259×10^4	-2.4205×10^4	-2.3998×10^4	-2.3900×10^4	-2.4770×10^4
RIO	-3.0736×10^4	-2.2551×10^4	-2.2014×10^4	-1.9302×10^4	-2.1230×10^4	-2.2184×10^4
NCP	-3.3026×10^4	-2.1101×10^4	-1.9878×10^4	-2.0052×10^4	-1.9405×10^4	-2.0778×10^4
WBC	-2.6584×10^4	-2.4763×10^4	-2.3715×10^4	-2.3932×10^4	-2.3697×10^4	-2.4334×10^4
FGL	-2.5552×10^4	-2.5128×10^4	-2.4451×10^4	-2.4848×10^4	-2.5070×10^4	-2.5022×10^4
CSR	-2.5797×10^4	-2.4363×10^4	-2.1567×10^4	-2.4120×10^4	-2.4045×10^4	-2.4080×10^4
BIL	-3.2613×10^4	-2.2380×10^4	-1.988×10^4	-2.1305×10^4	-1.8045×10^4	-2.1995×10^4
SGB	-3.4946×10^4	-2.3338×10^4	-2.2144×10^4	-2.1554×10^4	-1.4122×10^4	-2.3153×10^4

Table B.9: This table shows the log-likelihood values for the models estimated in Tables B.3 to B.8 using high frequency data with $\tau = 10$ minutes. The variance gamma of Madan and Seneta (1990) displays the largest (least negative) values.

St.Dev:	$1 < \sigma < 15$	$2 < \sigma < 15$	$3 < \sigma < 15$	$4 < \sigma < 15$	$5 < \sigma < 15$
Security					
AGL	2.9 (0.987)	3.2 (0.990)	3.4 (0.990)	3.6 (0.987)	3.2 (0.983)
ANZ	3.0 (0.989)	3.4 (0.994)	3.6 (0.995)	3.8 (0.995)	3.9 (0.994)
CBA	2.9 (0.992)	3.2 (0.997)	3.4 (0.997)	3.4 (0.994)	3.3 (0.990)
RIO	3.1 (0.983)	3.6 (0.989)	3.9 (0.990)	4.2 (0.989)	4.4 (0.985)
NCP	2.7 (0.984)	3.1 (0.989)	3.3 (0.987)	3.6 (0.987)	3.9 (0.982)
WBC	3.2 (0.986)	3.6 (0.992)	3.9 (0.993)	4.3 (0.995)	4.5 (0.993)
FGL	3.5 (0.977)	4.0 (0.982)	4.4 (0.982)	4.7 (0.980)	5.1 (0.976)
CSR	2.9 (0.988)	3.3 (0.992)	3.5 (0.991)	3.6 (0.987)	3.8 (0.986)
BIL	2.7 (0.989)	3.0 (0.994)	3.2 (0.994)	3.4 (0.995)	3.4 (0.993)
SGB	3.0 (0.990)	3.3 (0.994)	3.5 (0.994)	3.5 (0.987)	3.3 (0.984)
Av.Value:	3.0 ± 0.2	3.4 ± 0.2	3.6 ± 0.3	3.8 ± 0.4	3.9 ± 0.5

Table B.10: Estimates of the tail index, α_{L_t} , for L_t , taken over different ranges of standard deviation with corresponding values for Pearson's r in brackets. Average values for α_{L_t} were calculated over the total data set and are shown with 99% error bars.

Tail Index:	$\alpha_{ L_t }$	α_{V_t}	α_{N_t}
Security			
AGL	3.4 (0.990)	2.6 (0.997)	2.7 (0.999)
ANZ	3.6 (0.995)	3.9 (0.999)	3.3 (0.999)
CBA	3.4 (0.996)	3.8 (0.999)	3.2 (0.999)
RIO	3.9 (0.990)	3.7 (0.997)	4.4 (0.998)
NCP	3.3 (0.987)	3.6 (0.999)	2.9 (0.998)
WBC	3.9 (0.993)	2.9 (0.996)	4.0 (0.999)
FGL	4.4 (0.982)	2.8 (0.999)	2.7 (0.999)
CSR	3.5 (0.991)	2.8 (0.999)	4.0 (0.999)
BIL	3.2 (0.994)	2.1 (0.995)	2.9 (0.998)
SGB	3.5 (0.994)	2.5 (0.998)	3.8 (0.999)
Av. Value:	3.6 ± 0.3	3.1 ± 0.5	3.4 ± 0.5

Table B.11: Estimates of the tail indices for $|L_t|$, V_t and N_t calculated over the range of standard deviation $3 < \sigma < 15$ with corresponding values for Pearson's r in brackets. Average values were calculated over the total data set and are shown with 99% error bars.

	Hill Estimate			
	α_{L_t}	$\alpha_{ L_t }$	α_{V_t}	α_{N_t}
Security				
AGL	3.707	3.869	2.517	2.9117
ANZ	3.81	4.434	4.010	3.867
CBA	4.084	4.1171	3.639	3.7226
RIO	3.455	3.024	3.739	4.8885
NCP	2.454	2.825	3.339	3.189
WBC	3.789	4.072	3.053	4.116
FGL	3.661	4.7467	2.783	2.617
CSR	3.141	3.677	2.458	4.382
BIL	2.873	3.093	2.6177	3.4865
SGB	4.358	2.9515	2.491	4.179

Table B.12: Values for power law index α for L_t , $|L_t|$, V_t and N_t .

	ACF		
	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security			
AGL	0.543	0.361	0.541
ANZ	0.562	0.303	0.209
CBA	0.599	0.306	0.224
RIO	0.516	0.603	0.287
NCP	0.512	0.35	0.225
WBC	0.527	0.519	0.300
FGL	0.745	0.352	0.272
CSR	0.633	0.466	0.351
BIL	0.400	0.187	0.240
SGB	0.638	0.541	0.479

Table B.13: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained via the ACF.

	PSD			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	0.936	0.595	0.397	0.229
ANZ	1.071	0.509	0.464	0.330
CBA	0.940	0.541	0.453	0.283
RIO	0.940	0.707	0.536	0.350
NCP	0.936	0.520	0.378	0.264
WBC	0.923	0.591	0.564	0.317
FGL	0.977	0.572	0.465	0.325
CSR	1.010	0.560	0.500	0.365
BIL	0.904	0.612	0.380	0.223
SGB	1.050	0.680	0.568	0.390

Table B.14: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained via the PSD.

	Variance Plot			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	0.842	0.480	0.301	0.215
ANZ	1.001	0.525	0.326	0.215
CBA	0.900	0.533	0.392	0.280
RIO	0.926	0.497	0.506	0.302
NCP	0.967	0.395	0.340	0.308
WBC	0.950	0.515	0.454	0.246
FGL	0.931	0.501	0.424	0.330
CSR	1.031	0.506	0.502	0.301
BIL	1.030	0.430	0.484	0.255
SGB	0.923	0.521	0.419	0.371
Av. Value:	0.950 ± 0.048	0.490 ± 0.037	0.415 ± 0.060	0.282 ± 0.040

Table B.15: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained using variance plots.

	R/S Analysis			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	0.940	0.731	0.301	0.159
ANZ	0.900	0.464	0.345	0.281
CBA	0.830	0.567	0.376	0.206
RIO	0.934	0.722	0.403	0.285
NCP	0.944	0.563	0.289	0.247
WBC	0.960	0.506	0.455	0.235
FGL	0.967	0.593	0.331	0.1766
CSR	0.936	0.723	0.499	0.304
BIL	0.838	0.625	0.314	0.264
SGB	0.910	0.664	0.464	0.300

Table B.16: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained via rescaled range analysis.

	Wavelet LRD			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	0.986	0.6667	0.396	0.205
ANZ	0.990	0.600	0.320	0.180
CBA	0.870	0.559	0.264	0.165
RIO	1.017	0.661	0.291	0.227
NCP	0.903	0.501	0.336	0.325
WBC	0.939	0.552	0.426	0.159
FGL	0.949	0.652	0.259	0.250
CSR	0.988	0.626	0.552	0.230
BIL	0.800	0.642	0.404	0.363
SGB	1.050	0.618	0.382	0.400

Table B.17: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained through wavelet based methods.

	DEFA			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	1.046	0.691	0.453	0.182
ANZ	0.9500	0.5400	0.220	0.190
CBA	0.959	0.645	0.386	0.232
RIO	0.904	0.744	0.482	0.312
NCP	0.941	0.666	0.279	0.151
WBC	0.981	0.499	0.455	0.161
FGL	0.961	0.627	0.380	0.170
CSR	0.972	0.780	0.529	0.273
BIL	0.837	0.613	0.408	0.167
SGB	1.071	0.704	0.402	0.275

Table B.18: Values for long memory parameter β for L_t , $|L_t|$, V_t and N_t . Estimates were obtained using detrended fluctuation analysis.

	Scaling			
	β_{L_t}	$\beta_{ L_t }$	β_{V_t}	β_{N_t}
Security				
AGL	0.480	0.634	0.782	0.883
ANZ	0.467	0.641	0.789	0.863
CBA	0.503	0.646	0.818	0.891
RIO	0.528	0.647	0.745	0.843
NCP	0.531	0.760	0.853	0.919
WBC	0.478	0.651	0.664	0.846
FGL	0.467	0.615	0.732	0.828
CSR	0.465	0.627	0.618	0.787
BIL	0.568	0.659	0.897	0.903
SGB	0.417	0.634	0.695	0.824

Table B.19: Values for scaling parameter H for L_t , $|L_t|$, V_t and N_t . Estimates were obtained through variance behaviour.

ZRE Scaling	
H	
Security	
AGL	-1.287
ANZ	-1.568
CBA	-1.324
RIO	-1.048
NCP	-1.343
WBC	-1.439
FGL	-1.398
CSR	-1.320
BIL	-1.358
SGB	-1.294

Table B.20: Values for scaling parameter H for $\mathbf{P}[L_t = 0]$.

Security	Wavelet Scaling			
	H_{L_t}	$H_{ L_t }$	H_{V_t}	H_{N_t}
AGL	0.486	0.600	0.729	0.950
ANZ	0.564	0.731	0.833	0.905
CBA	0.572	0.732	0.811	0.968
RIO	0.485	0.600	0.871	0.978
NCP	0.495	0.700	0.818	0.865
WBC	0.556	0.709	0.759	0.920
FGL	0.522	0.679	0.894	0.930
CSR	0.507	0.642	0.696	0.894
BIL	0.562	0.629	0.753	0.838
SGB	0.520	0.621	0.847	0.842

Table B.21: Values for scaling parameter H for L_t , $|L_t|$, V_t and N_t . Estimates were obtained via wavelet methods.

Appendix C

Tables for Chapter 4

Parameters	$\tau = 10 \text{ min}$	$\tau = 1\text{hr}$
μ	0.8898	0.7104
σ	1.334	1.192
a'	-0.06153	1.532
Pearson's r :	0.9968 (0.8101)	0.9988 (0.9573)

Table C.1: Parameter values for μ , σ and a' for threshold Brownian motion fitted on small and large time intervals. The values of Pearson's r for the fitted models are shown in both cases. The values in brackets were obtained from the best fit for the normal distribution.

	Hypothesis Test	
	p -value	Reject
Security		
AGL	7.8×10^{-6}	YES
ANZ	1.3×10^{-4}	YES
CBA	7.8×10^{-4}	YES
RIO	0.4085	NO
NCP	5.2×10^{-7}	YES
WBC	0.0315	YES
FGL	0.047	YES
CSR	0.8627	NO
BIL	5.1×10^{-12}	YES
SGB	0.2425	NO

Table C.2: Hypothesis test of independence for Log-return and Waiting time. The hypothesis is rejected if the p -value is < 0.05 .

Security	β	Parameters		
		τ_0 (sec)	Pearson's r (Stretched)	Pearson's r (Exp.)
AGL	0.7626	82.64	0.9989	0.9666
ANZ	0.7652	43.72	0.9976	0.9563
CBA	0.7404	29.19	0.9998	0.9561
RIO	0.7193	61.11	0.9961	0.9595
NCP	0.6921	26.85	0.9973	0.9497
WBC	0.7736	48.75	0.9971	0.9722
FGL	0.7630	51.92	0.9984	0.9729
CSR	0.7962	106.99	0.9984	0.9836
BIL	0.6945	67.31	0.9959	0.9541
SGB	0.7029	128.89	0.9983	0.9839

Table C.3: Parameter values of β and the average waiting time, τ_0 for the empirical survival function. Pearson's r values are shown for the fitted stretched exponential (column (4)), and regular exponential (column (5)) distributions.

Appendix D

Programs for managing ASX data

This appendix contains a set of MATLAB functions that were written to create subsets of the ASX data set.

```
%%-----inputdata.m-----%  
%% This file contains %  
%% all the data %  
%% necessary to run the %  
%% program findinterval %  
%%-----%  
  
%% NOTE: For intra-day prices intsize must be set to 1,  
%% and for extra-day prices fracofhour must be set to 6.
```

```
function[usefile,inttype,fracofhour,intsize,initialtime,...  
        initialdate,finaltime,finaldate,stocks]=inputdata();
```

```
usefile      = 'yes';           % read data from this file?  
inttype      = 'min';          % 'min','hr','day' or 'week'  
fracofhour   = 10/60;          % intra-day interval size (hours)  
intsize      = 1;              % extra-day interval size (days)
```

```

initialtime = [10,00,00];
initialdate = [1,6,1993];
finaltime=[16,00,00];
finaldate   = [28,6,2002];

stocks = ['ANZ';'WBC';'AGL';'BIL';'CBA';'CSR';'FGL';'NCP';'RIO'];

disp(fracofhour),disp('Stocks:'),disp('-----'),disp(stocks)

%%-----findinterval.m-----%
%% This program extracts a subset of data      %
%% from the Stock directory. The extracted set %
%% should consist of prices over a specified  %
%% interval.                                   %
%%-----%

% =====
% Create vector of all times and dates needed
% =====

clear

[itimedate,ftimedate,dt,fracofhour,...
stocklist,intervaltype,intervalsize]= startstoptimes;

tenam      = time2num([10,0,0.30],[0,0,0]);
fourpm     = time2num([16,0,0.30],[0,0,0]);

timesneeded=tvector(itimedate,ftimedate,dt,...
fracofhour,intervaltype,intervalsize);

    disp('Checkpoint 1')
yy=datevec(timesneeded);
tsttens=datenum(0,0,0,yy(:,4),yy(:,5),yy(:,6));

```

```

tens=find(tsttens==time2num([10,0,0],[0,0,0])); tsttens=[];

% -----
% Filter data file with vector of desired times
% =====
for jj=1:length(stocklist);
    stockname=stocklist(jj)
    [price,volume,time,cap]=readsirca('pt',stockname);    % load data

    disp('Checkpoint 2')
    ti=length(time);
    if mod(ti,2)==0
        yr1=datevec(time(1:ti/2));
        yr2=datevec(time((ti/2)+1:ti));
        yr=cat(1,yr1,yr2);
        yr1=[];
        yr2=[];
    else
        yr1=datevec(time(1:(ti+1)/2));
        yr2=datevec(time(((ti+1)/2)+1:ti));
        yr=cat(1,yr1,yr2);
        yr1=[];
        yr2=[];
    end

    tsttime=datenum(0,0,0,yr(:,4),yr(:,5),yr(:,6));
    yr=[];

    % remove erroneous entries from data set
    %-----
    psi=find((tsttime<=tenam)|(tsttime>fourpm));
    time(psi,:) = [];
    price(psi) = [];

```

```

%volume(psi)= [];
%volume=cumsum(volume);      % Cumulatively sum volume

tsttime=[];

sampletime = zeros(length(timesneeded),1);
sampleprice = sampletime;
%samplevol  = sampletime;

disp('Checkpoint 3')

% match required times with corresponding data set times
%-----
opentime=[0,0,0];
tout=cputime;
for ii=1:length(timesneeded)
    if ~ismember(ii,tens)
        dj      = max(find(time<timesneeded(ii)));
% Make sure the time is greater than that of the opening price.
        if time(dj)<opentime(2)
            sampletime(ii) = opentime(2);
            sampleprice(ii)= opentime(1);
            %samplevol(ii)  = opentime(3);
            time(1:dj-1)=[];
            price(1:dj-1)=[];
            %volume(1:dj-1)=[];
        else
            sampletime(ii) = time(dj);
            sampleprice(ii)= price(dj);
            %samplevol(ii)  = volume(dj);
            time(1:dj-1)=[];
            price(1:dj-1)=[];
            %volume(1:dj-1)=[];
        end
    end
end

```

```

        end
    else
        dj      = min(find(time>timesneeded(ii)));
        samptime(ii) = time(dj);
        sampleprice(ii)= price(dj);
        %samplevol(ii) = volume(dj);
        opentime=[sampleprice(ii),samptime(ii)];
    end

end

time=[];psi=[];tsttime=[];price=[]; % Clear for next loop
tend=cputime-tout
disp('Checkpoint 4')

% Write data to file
%-----
switch(intervaltype)
    case {'hr','min'}
        if strcmp(intervaltype,'min')
            filename=cat(2,char(stockname),...
                num2str(60*fracofhour),char(intervaltype),'.dat');
        else
            filename=cat(2,char(stockname),...
                num2str(fracofhour),char(intervaltype),'.dat');
        end
    case 'day'
        filename=cat(2,char(stockname),...
            num2str(intervalsize),char(intervaltype),'.dat');
    case 'week'
        filename=cat(2,char(stockname),...
            num2str(intervalsize),char(intervaltype),'.dat');
end
end

```

```

    fid = fopen(filename,'w');
    fprintf(fid,'%f %f %f \n',[sampleprice';sampletime';timesneeded']);
    fclose(fid);
end

%%-----readsirca.m-----%
%% This function reads data from the data file %
%% *.csv. The user inputs the desired      %
%% data columns (eg. s(stockprice), v(olume), %
%% t(ime and date), c(apitalisation)) and the %
%% program reads the data into a matrix b.   %
%%-----%

function [price,volume,timeday,cap]=readsirca(datainput,stockname)
%datainput=input('What data is to be stored? (pvt dc): ','s');
possiblechoice=['p','v','t','c'];
long=zeros(1,4);      % Specifies how many columns for each of pvt dc
price=[];volume=[];timeday=[];cap=[];

filename=cat(2,char(stockname),'.csv');
fid = fopen(filename);
y = ['%g '];          % read number
n = ['%*g'];          % ignore number
output=char(n,n,n,n,n);
rcl=char(n,n,n,n,n,n,n,n,n);
for ii=1:4
    if ~isempty(find(datainput==possiblechoice(ii)))
        ~ switch(possiblechoice(ii))
            case 'p'    % Price
                rcl(1,1:3)=y;
                long(1)=1;
                output(1,1:3)=y;
            case 'v'    % Volume

```

```

        rcl(2,1:3)=y;
        long(2)=1;
        output(2,1:3)=y;
    case 't'    % Time and Date
        rcl(3,1:3)=y;
        rcl(4,1:3)=y;
        rcl(5,1:3)=y;
                rcl(6,1:3)=y;
        rcl(7,1:3)=y;
        rcl(8,1:3)=y;
        long(3)=6;
        output(3,1:3)=y;
    case 'c'    % Capitalisation
        rcl(9,1:3)=y;
        long(4)=1;
        output(4,1:3)=y;
    otherwise
        warning('Warning: non-standard input!!')
    end
end
end
end
numberofcolumns=cumsum(long);
aa= [rcl(1,:) ',' rcl(2,:) ',' rcl(3,:) ':' rcl(4,:) ,...
     ':' rcl(5,:) ',' rcl(6,:) '/' rcl(7,:) '/' rcl(8,:) ',' rcl(9,:) ];
a = fscanf(fid,aa,[numberofcolumns(end) inf]);
fclose(fid);
for jj=1:5
    if output(jj,1:3)==y
        if jj==1
            price = a(1,:);
            a(1,:)=[];
        elseif jj==2
            volume = a(1,:);

```

```

        a(1,:)=[];
    elseif jj==3
        td = a(1:long(jj),:);
    a(1:long(jj),:)=[];
        timeday=datetime(td(:,6),td(:,5),td(:,4),td(:,1),td(:,2),td(:,3));
    td=[];
    elseif jj==4
        cap = a(1,:);
        a(1,:)=[];
    end
end
end
end

%%-----readlist.m-----%
%% This function reads from the data %
%% contents file stock.list and %
%% returns the list of stocks, first %
%% times recorded and rank in top 200.%
%%-----%

function [tickname,timeday,score]=readlist(stocks)

[tickname,hr,mn,sc,dy,mth,yr,score]...
    =textread('stock.list','%[^\n],%d:%d:%d,%d/%d/%d,%*f,%d');

timeday=datetime([yr,mth,dy,hr,mn,sc]);
hr=[];mn=[];sc=[];dy=[];mth=[];yr=[];

if isempty(stocks)
    stocklist=input('Stocks wanted (List, Rank or Top XX): ');
else
    stocklist=stocks;
end
end

```

```

% Case when we want certain stocks
%=====
if ischar(stocklist)
    listofwantedstocks=stocklist;
    xi=find(~ismember(tickname,listofwantedstocks));
    tickname(xi) = [];
    timeday(xi) = [];
    score(xi) = [];

% Case when we want the top XX stocks
%=====
elseif ~ischar(stocklist)&(length(stocklist)==1)
    topnumber=stocklist;
    xi=find(score>topnumber);
    tickname(xi) = [];
    timeday(xi) = [];
    score(xi) = [];
    if isempty(tickname)
        error('Only Top 199 stocks available!!!')
    end

% Case when we want the stocks specified by rank
%=====
elseif ~ischar(stocklist)&(length(stocklist)>1)
    listofnumbers=stocklist;
    xi=find(~ismember(score,listofnumbers));
    tickname(xi) = [];
    timeday(xi) = [];
    score(xi) = [];
    if length(tickname)~=length(listofnumbers)
        error('Only Top 199 stocks available!!!')
    end
else

```

```

    error('Error: Unrecognised input!')
end

%%-----startstoptimes.m-----%
%% This function retrieves initial and final %
%% times and date for findinterval.m. It also %
%% returns the list of stocks wanted.      %
%%-----%

function [itimedate,ftimedate,dt,...
fracofhour,stocklist,inttype,intsize]=startstoptimes()

[usefile,inttype,fracofhour,intsize,...
initialtime,initialdate,finaltime,finaldate,stocks]=inputdata;

switch(usefile)
    case 'no'
        stocks=[];
        inttype = input('Enter interval type (day,week): ','s');
        switch(inttype)
            case {'hr','min'}
                fracofhour = input('Enter intra-day interval (in hours): ');
                if mod(6,fracofhour)~=0|fracofhour>6
                    error('Error: intra-day interval must be factor of 6 <= 6')
                end
                intsize = 1;
            case 'day'
                fracofhour = 6;
                intsize = input('Enter daily interval (in days): ');
            case 'week'
                fracofhour = 6;
                intsize = input('Enter weekly interval (in weeks): ');
        end
end

```

```

    initialtime = input('Initial time ([HH,MM,SEC]): ');
    initialdate = input('Initial date ([dd,mm,yyyy]): ');
    finaltime   = input('Final time ([HH,MM,SEC]): ');
    finaldate   = input('Final date ([dd,mm,yyyy]): ');
    case 'yes'
    otherwise
        error('unknown command')
    end

itimedate = time2num(initialtime,initialdate);
ftimedate = time2num(finaltime,finaldate);
dt        = time2num([fracofhour,0,0],[0,0,0]);

% check initial and final dates for non-trading days
%=====
if whatdayvec(initialdate(1),initialdate(2),initialdate(3))==0
    fprintf('\n Warning: Initial date is a Sunday, please re-enter. \n\n');
    return
elseif whatdayvec(initialdate(1),initialdate(2),initialdate(3))==6
    fprintf('\n Warning: Initial date is a Saturday, please re-enter. \n\n');
    return
elseif whatdayvec(finaldate(1),finaldate(2),finaldate(3))==0
    fprintf('\n Warning: Final date is a Sunday, please re-enter. \n\n');
    return
elseif whatdayvec(finaldate(1),finaldate(2),finaldate(3))==6
    fprintf('\n Warning: Final date is a Saturday, please re-enter. \n\n');
    return
end

% check files to be read and make sure itimedate and
% ftimedate are subsets of all data files. If not, display error:
% Warning: First date in "XXX.csv, No. XXX" beginning at "xx:xx:xx xx/xx/xx"
%=====

```

```
[stocklist,starttimes,score]=readlist(stocks);

xi=find(starttimes>=itimedate);
if ~isempty(xi)
    psi=find(starttimes==max(starttimes(xi)));
    fprintf('\n Warning: First date in %s.csv, No.%d beginning at %s \n\n',...
        char(stocklist(psi)),score(psi),datestr(starttimes(psi)));
    clear
    return
end
```

Appendix E

Program for calculating the large deviation function

This appendix contains a FORTRAN program that solves Eq. (4.109) and calculates the large deviation function $\Phi(u)$.

```
C C    PROGRAM OrnsteinLDF C
      IMPLICIT NONE
      INTEGER :: i, nint
      REAL (KIND=8) :: a, r, dr, PhiR
C
      OPEN(UNIT=15,FILE='PhiR.DAT')
      WRITE(*,*) ' Value for a and number of intervals: '
      READ(*,*) a, nint
C
      IF(a .LT. 1.D-2) a=1.D-2
      dr = 1.D0/REAL(nint-1)
      r = -dr
      DO i = 1,nint
         r = r+dr
         WRITE(15,*) r, PhiR(r, a)
      ENDDO
      CLOSE (15)
```

APPENDIX E. PROGRAM FOR CALCULATING THE LARGE DEVIATION FUNCTION 226

```

        END
C C
*****
C
    REAL (KIND=8) FUNCTION PhiR(r, a)
    IMPLICIT NONE
    REAL (KIND=8) :: r, a
    REAL (KIND=8), PARAMETER :: eps=1.D-4
    REAL (KIND=8) :: smin, smax
    REAL (KIND=8) :: NegFn, val
    INTEGER, PARAMETER :: max=100
    INTEGER :: i, icon
    REAL (KIND=8) :: ss, aa, ar2, rr
    REAL (KIND=8) :: s0, s1, ds, xminbrent
    EXTERNAL NegFn
C
    COMMON /funcs/ ss, aa, ar2, rr
C
    aa = a
    rr = r
    smin = -10.D0
    smax = 10.D0
    ds = 0.5D0
C C
-----
    s0 = xminbrent(smin,smax,NegFn,eps)
    PhiR = -NegFn(s0)
C
-----
    RETURN
    END
C C
*****

```

APPENDIX E. PROGRAM FOR CALCULATING THE LARGE DEVIATION FUNCTION 227

```

C
  REAL (KIND=8) FUNCTION NegFn(s)
  IMPLICIT NONE
  REAL (KIND=8) :: s
  REAL (KIND=8) :: ss, a, ar2, rr, ESolve
  INTEGER :: m, iw

C
  COMMON /funcs/ ss, a, ar2, rr

C
  NegFn = rr*s+0.5- ESolve(s)

C
  RETURN
  END

C C
*****
C
  REAL (KIND=8) FUNCTION ESolve(s)
  IMPLICIT NONE
  REAL (KIND=8), PARAMETER :: eps=1.D-6, root2=1.414213562373095D0
  REAL (KIND=8) :: e, s, a, x, ss, ar2, x0, x1, rr
  REAL (KIND=8) :: sefun, dxfun, rt1, rt2
  INTEGER :: i, icon
  EXTERNAL sefun, dxfun

C
  COMMON /funcs/ ss, a, ar2, rr
C C Find an upperbound for the eigenvalue
  ss = s
  x0 = 0.D0
  ar2 = -root2*a
  x1 = 0.9999D0

C
  CALL DTSD1(x0, x1, dxfun, eps, rt1, icon)
  x0 = 0.D0

```

APPENDIX E. PROGRAM FOR CALCULATING THE LARGE DEVIATION FUNCTION 228

```

    ar2 = -ar2
    x1 = 1.D0+a+a
    CALL DTSD1(x0, x1, dxfun, eps, rt2, icon)
    rt2 = rt2+s
    x1 = MIN(rt1,rt2)-eps
C C Find a lower bound for the eigenvalue
    x0 = x1-ABS(s)
    rt1 = sefun(x1)
    DO WHILE(sefun(x0)*rt1 .GT. 0.D0)
        x0 = x0-0.1
    ENDDO
C C Solve for the eigenvalue
    ar2 = ABS(ar2)
    CALL DTSD1(x0, x1, sefun, eps, x, icon)
    ESolve = x+0.5D0
C
    RETURN
    END
C C
*****
C
    REAL (KIND=8) FUNCTION sefun(x)
    IMPLICIT NONE
    REAL (KIND=8) :: x, ss, a, ar2, rr, s1, s2
    REAL (KIND=8) :: dxfun
c
    COMMON /funcs/ ss, a, ar2, rr
C
    ar2 = ABS(ar2)
    s1 = (x-ss)*dxfun(x-ss-1.D0)/dxfun(x-ss)
    ar2 = -ar2
    s2 = x*dxfun(x-1.D0)/dxfun(x)
    sefun = s1+s2

```

APPENDIX E. PROGRAM FOR CALCULATING THE LARGE DEVIATION FUNCTION 229

```
      RETURN
      END
C C
*****
C
      REAL (KIND=8) FUNCTION dxfun(x)
C C This fuction calculates parabolic cylinder functions C using
the subtourine pbdv C
      IMPLICIT NONE
      REAL (KIND=8) :: x, ss, a, ar2, rr
      REAL (KIND=8) :: pdf, pdd
      REAL (KIND=8) :: dv(0:100),dp(0:100)
c
      COMMON /funcs/ ss, a, ar2, rr
C
      CALL pbdv(x, ar2, dv, dp, pdf, pdd)
      dxfun = pdf
      RETURN
      END
C C
*****
```

Bibliography

- [1] Abhyankar, A., Ghosh, B., Levin, E. and Limmack, R. (1997). Bid-ask spreads, trading volume and volatility: Intra-day evidence from the London stock exchange. *Journal of Business Finance & Accounting*, 24, 343-362 .
- [2] Abramowitz, M. and Stegun, I. (1968). *Handbook of Mathematical Functions*. Dover, New York.
- [3] Abry, P., Flandrin, P., Taqqu, M. and Veitch, D. (2000). Wavelets for the analysis estimation and synthesis of scaling data. *Self-Similar Network Traffic and Performance Evaluation*, Wiley, New York, 39-88.
- [4] Abry, P. Gonçalvés, P. and Flandrin, P. (1995). Wavelets, spectrum estimation and 1/f processes. *Wavelets and Statistics*, Lecture Notes in Statistics, 103, Springer-Verlag, 15-30.
- [5] Alexander, S.S. (1961). Price movements in speculative markets: trends or random walks. *Industrial Management Review*, 2, 7-26.
- [6] Amaral, L.A.N., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., Stanley, H.E. and Stanley, M.H.R. (1997). Scaling behaviour in economics: I. Empirical results for company growth. *Journal de Physique I France*, 7, 621-633.
- [7] Andersen, T. and Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4, 115-158.

- [8] Antoniou, I., Ivanov, Vi., Ivanov, Va. and Zrelov, P. (2004). On the log-normal distribution of stock market data. *Physica A*, 331, 617-638.
- [9] Barndorff-Nielsen, O.E. (1995). Normal inverse gaussian and the modelling of stock returns. *Research Report*, Aarhus University.
- [10] Benston, G.J. and Hagerman, R.L. (1974). Determinants of bid-asked spreads in the over-the-counter market. *Journal of Financial Economics*, 1, 353-364.
- [11] Beran, J. (1994). *Statistics for Long-Memory Processes*. Chapman & Hall/CRC, London.
- [12] Bertram, W.K. and Peiris, M.S. (2004). An example of a misclassification problem applied to Australian equity data. *Submitted to Economics Letters*.
- [13] Bertram, W.K. (2004). An empirical investigation of Australian Stock Exchange data. *Physica A*, 341, 533-546.
- [14] Bertram, W.K. (2005). A threshold model for Australian equities, *Physica A*, 346, 561-577.
- [15] Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637-654.
- [16] Blattberg, R.C. and Gonendes, N.J. (1974). A comparison of the stable and student distributions as statistical models for stock prices. *Journal of Business*, 47, 244-280.
- [17] Blattberg, R.C. and Gonendes, N.J. (1977). A comparison of the stable and student distributions as statistical models for stock prices: Reply. *Journal of Business*, 50, 78-79.
- [18] Bochner, S. (1955). *Harmonic analysis and the theory of probability*. University of California Press, Berkeley.
- [19] Bollerslev, T. (1986). Generalised Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31, 307-327.

- [20] Bollerslev, T. and Ghysels, E. (1996). Periodic Autoregressive Conditional Heteroscedasticity. *Journal of Business and Economic Statistics*, 14, 139-151.
- [21] Borodin, A. and Salminen, P. (2002). *Handbook of Brownian Motion*. Birkhäuser Verlag, Berlin.
- [22] Bouchaud, J. and Potters, M. (1997). *Theories des Risques Financiers*. Eyrolles, Alea-Saclay, 1997.
- [23] Bouchaud, J. Potters, M. and Meyer, M. (2000). Apparent multifractality in financial time series. *The European Physical Journal B*, 13, 595-599.
- [24] Box, G.E.P. and Jenkins, G.M. (1970). *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco.
- [25] Buchen, P. and Mainardi, F. (1975). Asymptotic expansions for transient viscoelastic waves. *Journal de Mécanique*, 14, 597-608.
- [26] Buldyrev, S.V., Amaral, L.A.N., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., Stanley, H.E. and Stanley, M.H.R. (1997). Scaling behaviour in economics: II. Empirical results for company growth. *Journal de Physique I France*, 7, 635.
- [27] Calvert, L., Fisher, A., and Mandelbrot, B.B. (1997). Large Deviations and the Distribution of Price Change. *Cowles Foundation Discussion Paper #1166*.
- [28] Caputo, M. and Mainardi, F. (1971). Linear models of dissipation in anelastic solids. *La Rivista del Nuovo cimento*, 1, 161-198.
- [29] Carr, P., Geman, H., Madan, D. and Yor, M. (2002). The fine structure of asset returns: an empirical investigation. *Journal of Business*, 75, 305-332.
- [30] Clark, P.K. (1973). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica*, 41, 135-159.
- [31] Cootner, P.H. (Ed.) (1964). *The random character of the stock market*. MIT press, Cambridge.

- [32] Costa, R.L. and Vasconcelos, G.L. (2003). Long-range correlations and non-stationarity in the Brazilian stock market. *Physica A*, 329, 231-248 .
- [33] Cox, J.C. and Ross, S.A. (1976). The valuation of options for alternative stochastic processes, *Journal of Financial Economics*, 3, 145-166.
- [34] Dacorogna, M.M., Gencay, R., Muller, U.A., Olsen, R.B. and Pictet, O.V. (2001). *An Introduction to High Frequency Finance*. Academic Press, London.
- [35] Dekkers, A. and de Haan, L. (1993). Optimal Choice of sample fraction in extreme-value estimation. *Journal of multivariate analysis*, 47, 173-195.
- [36] Ding, Z. and Granger, C.W.J. (1996). Modelling volatility persistence of speculative returns: a new approach. *Journal of Econometrics*, 73, 185-215.
- [37] Duan, J.C. (1995). The GARCH option pricing model. *Mathematical Finance*, 5, 15-32.
- [38] Eberlien, E. and Keller, K. (1995). Hyperbolic distributions in finance. *Bernoulli*, 1, 281-299.
- [39] Edelman, D. and Gillespie, T. (1997). The stochastically subordinated log-normal process applied to financial time series and option pricing. *Research Report*, University of Sydney.
- [40] Engle, R.F. (1982). Auto-Regressive Conditional Heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 50, 987-1008.
- [41] Fama, E.F. (1963). Mandelbrot and the Paretian hypothesis. *Journal of Business*, 38, 34-105.
- [42] Fama, E.F. (1965). The Behaviour of Stock-Market Prices. *Journal of Business*, 38, 34-105.
- [43] Feller, W. (1966). *Introduction to Probability*. Volume II. Wiley, New York.

- [44] Fisher, A., Calvert, L. and Mandelbrot, B.B. (1997). Multifractality of Deutschmark/US Dollar Exchange Rates. *Cowles Foundation Discussion Paper* #1165.
- [45] Fusai, G. (2000). Corridor options and arc-sine law. *The Annals of Applied Probability*, 10, 634-663.
- [46] Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. (2003a). A theory of power-law distributions in financial market fluctuations, *Nature*, Vol 423 , 267-270.
- [47] Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. (2003b). Understanding the cubic and half-cubic laws of financial fluctuations. *Physica A*, 324, 1-5.
- [48] Geman, H. and Ane, T. (1996). Stochastic Subordination. *Risk*, 9, 146-149.
- [49] Gopikrishnan, P., Meyer, M., Amaral, L.A.N., Plerou, V. and Stanley, H.E. (1998). Inverse Cubic Law for distribution of Stock Price Variations. *The European Physical Journal B*, 3, 139-140.
- [50] Gopikrishnan, P., Plerou, V., Amaral, L.A.N., Meyer, M. and Stanley, H.E. (1999). Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60 , 5305-5316.
- [51] Gopikrishnan, P., Plerou, V., Liu, Y., Amaral, L.A.N., Gabaix, X. and Stanley, H.E. (2000). Scaling and correlation in financial time series. *Physica A*, 287, 362-373.
- [52] Gorski, A.Z., Drozd, S. and Speth, J. (2002). Financial multifractality and its subtleties: an example of DAX. *Physica A*, 316, 496-510.
- [53] Granger, C.W.J. (1966). The typical spectral shape of an economic variable. *Econometrica*, 34, 150-161.
- [54] Granger, C.W.J. and Ding, Z. (1996). Varieties of long memory models. *Journal of Econometrics*, 73, 61-77.

- [55] Granger, C.W.J. and Teräsvirta, T. (1999). A simple nonlinear time series model with misleading linear properties. *Economics Letters*, 62, 161-165.
- [56] Feder, J. (1988). *Fractals*. Plenum Press, New York.
- [57] Hall, P. (1982). On simple estimates of an exponent of regular variation. *Journal of the Royal Statistical Society, Series B*, 44, 37-42.
- [58] Hastings, N. and Peacock, J. (1974). *Statistical Distributions*. Butterworth and Co. Ltd., London.
- [59] Heston, S. (1993). A closed form solution for options with stochastic volatilities with applications to bond and currency options. *Review of Financial Studies*, 6, 327-343.
- [60] Heston, S. and Nandi, S. (1997). A closed form GARCH option pricing model. *Federal Reserve Bank of Atlanta, Working paper 97-9*.
- [61] Heyde, C.C. (1999). A minimal description model for asset prices. *Journal of Applied Probability*, 36, 1234-1239.
- [62] Heyde, C.C. and Gay, R. (2000). Fractals and contingent claims. *Preprint*.
- [63] Heyde, C.C., and Liu, S. (2001). Empirical realities for a minimal description risky asset model. *J. Korean Math. Soc.*, 38, 1047-1059.
- [64] Hilfer, R. (1994). *Stochastische Modelle für die betriebliche Planung*. GBI-Verlag, Munich.
- [65] Ho, D., Lee, C., Wang, C. and Chuang, M. (2004). Scaling characteristics in the Taiwan stock market. *Physica A*, 332, 448-460.
- [66] Hughes, B. (1995). *Random Walks and Random Environments, Vol. 1: Random Walks*. Clarendon Press, Oxford.
- [67] Hull, J.C. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42, 281-300.

- [68] Hurst, H. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116, 770-799.
- [69] Hurst, S., Platen, E. and Rachev, S. (1997). Subordinated Market Index Models. *Financial Engineering and the Japanese Markets*, 4, 97-124.
- [70] Itô, K. (1951). On stochastic differential equations. *Memoirs, American Mathematics Society*, 4, 1-51.
- [71] Jolliffe, I.T. (1986). *Principal Component Analysis*. Springer, New York.
- [72] Kac, M. (1949). On Distributions of Certain Wiener Functionals. *Transactions of the American Mathematical Society*, 65, 1-13.
- [73] Kac, M. and Darling, D. (1957). On Occupation Times for Markoff Processes. *Transactions of the American Mathematical Society*, 84, 444-458.
- [74] Kallsen, J. and Taqqu, M. (1998). Option pricing in ARCH-type models. *Mathematical Finance*, 8, 13-26.
- [75] Kamke, E. (1959). *Differentialgleichungen loesungsmethoden und loesungen*. Chelsea Publishing Company, New York.
- [76] Karatzas, I. and Shreve, S.E. (1988). *Brownian motion and stochastic calculus*. Springer, Berlin.
- [77] Kendall, M.G. (1953). The analysis of economic time series. *Journal of the Royal Statistical Society*, 96, 11-25.
- [78] Khintchine, A.Ya. and Levy, P. (1936). Sur les loi stables. *C.R. Acad. Sci. Paris*, 202, 374-376.
- [79] Kruiženga, R. (1956). *Put and Call options: a theoretical and market analysis*. PhD dissertation, MIT.
- [80] Lamperti, J. (1958). An Occupation Time Theorem for A Class of Stochastic Processes. *Transactions of the American Mathematical Society*, 88, 380-387.

- [81] Levy, P. (1925). *Calcul des probabilités*. Gauthier-Villars, Paris.
- [82] Levy, P. (1939). Sur certains processus stochastique homogènes. *Compositio Mathematica*, 7, 283.
- [83] Lewis, A. (2000). *Option valuation under stochastic volatility*. Finance Press, California.
- [84] Luckock, H. (2003). A steady state model of the continuous double auction. *Quantitative Finance*, 3, 385-404.
- [85] Madan, D. and Seneta, E. (1987a). Simulation of estimates using the empirical characteristic function. *International Statistical Review*, 55, 153-161.
- [86] Madan, D. and Seneta, E. (1990). The variance gamma model for share market returns. *Journal of Business*, 63, 511-524.
- [87] Madan, D. (2002). Purely discontinuous asset price processes. *Handbook of Mathematical Finance*. Cambridge University Press, Cambridge, 105-153.
- [88] Mainardi, F., Raberto, M., Gorenflo, R. and Scalas, E. (2000). Fractional calculus and continuous time finance II: the waiting time distribution. *Physica A*, 287, 468-481.
- [89] Majumdar, S. and Bray, A. (2002). Large-deviation functions for nonlinear functionals of a Gaussian stationary Markov process. *Physical Review E*, 65, 051112.
- [90] Mandelbrot, B.B. (1959). Variables et processus stochastiques de Pareto-Levy, et la repartition des revenus. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, 23, 2153-2155.
- [91] Mandelbrot, B.B. (1960). The Pareto-Levy law and the distribution of income. *International Economic Review*, 1, 79-106.
- [92] Mandelbrot, B.B. (1963). The variation of certain speculative prices. *Journal of Business*, 36, 394-419.

- [93] Mandelbrot, B.B. and Taylor, H.M. (1967). On the Distribution of Stock Price Differences, *Operations Research*, 15, 1057-1062.
- [94] Mandelbrot, B.B. and Wallis, J.R. (1968). Robustness of the rescaled range R/S and the measurement of non cyclic long-run statistical dependence. *Water Resources Research*, 4, 909-918.
- [95] Mandelbrot, B.B., Fisher, A. and Calvert, L. (1997). A Multifractal Model of Asset Returns. *Cowles Foundation Discussion Paper #1164*.
- [96] Mantegna, R.N. and Stanley, H.E. (1994). Stochastic process with ultraslow convergence to a gaussian: the truncated Levy flight. *Phys.Rev. Lett.*, 73, 2946-2949.
- [97] Mantegna, R.N. and Stanley, H.E. (1995). Scaling behaviour in the dynamics of an economic index. *Nature*, 383, 587-588.
- [98] Mantegna, R.N. and Stanley, H.E. (1999). *Introduction to Econophysics: Correlation and Complexity in Finance*. Cambridge University Press, Cambridge.
- [99] Martens, M., Chang, Y. and Taylor, S. (2002). A comparison of seasonal adjustment methods when forecasting intraday volatility. *Journal of Financial Research*, 25, 283-299.
- [100] Merton, R.C. (1976). Option pricing when the underlying stock returns are discontinuous, *Journal of Financial Economics*, 5, 125-144.
- [101] Mittnik, S. and Rachev, S. (1993a). Modelling asset returns with alternative stable distributions. *Econometric Reviews*, 12, 261-330.
- [102] Mizuno, T., Kurihara, S., Takayasu, M. and Takayasu, H. (2003). Analysis of high-resolution foreign exchange data of USD-JPY for 13 years. *Physica A*, 324, 296-302.
- [103] Molgedey, L. and Ebeling, W. (2000). Intraday patterns and local predictability of high-frequency financial time series. *Physica A*, 287, 420-428.

- [104] Montroll, E. and Weiss, G. (1967). Random Walks on lattices, II. *Journal of Mathematical Physics*, 6, 167-181.
- [105] Moore, A.B. (1962). *A statistical analysis of common stock prices*. PhD dissertation, University of Chicago.
- [106] Musiela, M. and Rutkowski, M. (1997). *Martingale methods in financial modelling*. Springer, Berlin.
- [107] Osbourne, M.F.M. (1959). Brownian motion in the stock market. *Operations Research*, 7, 145-173.
- [108] Pareto, V. (1896). *Cours d'economie politique*. Reprinted in Pareto, V. (1965), *oeuvres completes*, Geneva.
- [109] Peiris, M.S. (2004). Improving the Quality of Forecasting using Generalised AR Models: An Application to Statistical Quality Control. *Statistical Methods*, 5(2), 156-171.
- [110] Peiris, M.S., Allen, D. and Thanvaneswaran, A. (2004). An introduction to Generalized Moving Average Models and Applications. *Journal of Applied Statistical Science* (to appear).
- [111] Peng, C., Buldyrev, S., Havlin, S., Simons, M., Stanley, H.E. and Goldberger, A. (1994). Mosaic organization of DNA nucleotides. *Physical Review E*, 49, 1685-1689.
- [112] Praetz, P.D. (1969). Australian share prices and the random walk hypothesis. *Australian Journal of Statistics*, 11, 123-139.
- [113] Praetz, P.D. (1972). The distribution of share price changes. *Journal of Business*, 45, 49-55.
- [114] Praetz, P.D. (1977). A comparison of the stable and student distributions as statistical models for stock prices: Comment. *Journal of Business*, 50, 76-77.

- [115] Press, S.J. (1967). A compound events model for security prices. *Journal of Business*, 40, 317-335.
- [116] Raberto, M., Scalas, E. and Mainardi, F. (2002). Waiting times and returns in high frequency financial data: an empirical study. *Physica A*, 314, 749-755.
- [117] Rachev, S. and Mittnik, S. (2000). *Stable Paretian models in finance*. Wiley, New York.
- [118] Riedi, R.H. (2003). *Multifractal Processes. Theory and applications of Long-Range Dependence*. Birkhauser, Boston.
- [119] Renyi, A. (1970). *Probability Theory*. North-Holland, Amsterdam.
- [120] Rydberg, T. and Shepard N. (1999). Modelling trade-by-trade price movements of multiple assets using multivariate compound Poisson processes. Nuffield College, Oxford, *Working paper series* 1998-W19.
- [121] Samuelson, P.A. (1955). Brownian motion in the stock market. *Unpublished Manuscript*.
- [122] Samuelson, P.A. (1965). Rational theory of warrant pricing. *Industrial management review*, 6, 13-32.
- [123] Scalas, E., Gorenflo, R. and Mainardi, F. (2000). Fractional calculus and continuous time finance. *Physica A*, 284, 376-384.
- [124] Scalas, E., Gorenflo, R., Luckock, H., Mainardi, F., Mantelli, M. and Raberto, M. (2004). Anomalous waiting times in high frequency financial data. *preprint*.
- [125] Scher, H. and Lax, M. (1973). Stochastic transport in a disordered solid. I. Theory. *Physical Review*, 137, 4502-4519.
- [126] Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, London.
- [127] Stanislavsky, A.A. (2003). Black-Scholes model under subordination. *Physica A*, 318, 469-474.

- [128] Stein, E. and Stein, J. (1991). Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies*, 4, 727-752.
- [129] Taylor, S. and Xu, X. (1997). The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance*, 4, 317-40.
- [130] Toroczkai, Z., Newman, T. and Das Sarma, S. (1999). *Physical Review E*, 60, R1115.
- [131] Weiss, G. and Calabrese, P. (1996). Occupation times of a CTRW on a lattice with anomalous sites. *Physica A*, 234, 443-454.
- [132] Wiener, N. (1923). Differentiable Space. *Journal of Mathematics and Physics*, 2, 131.
- [133] Wood, R., McInish, T. and Ord, A. (1985). An Investigation of Transactions Data for NYSE Stocks. *The Journal of Finance*, 40, 723-739.

- 1 JUL 2005

UNIVERSITY OF SYDNEY LIBRARY

0000000606265369

RARE BOOKS LIB