

Cambridge University Press

0521770416 - Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses and Dick Van Dijk

Excerpt

[More information](#)

# 1 Introduction

---

This book deals with the empirical analysis of financial time series with an explicit focus on, first, describing the data in order to obtain insights into their dynamic patterns and, second, out-of-sample forecasting. We restrict attention to modelling and forecasting the conditional mean and the conditional variance of such series – or, in other words, the return and risk of financial assets. As documented in detail below, financial time series display typical nonlinear characteristics. Important examples of those features are the occasional presence of (sequences of) aberrant observations and the plausible existence of regimes within which returns and volatility display different dynamic behaviour. We therefore choose to consider only nonlinear models in substantial detail, in contrast to Mills (1999), where linear models are also considered. Financial theory does not provide many motivations for nonlinear models, but we believe that the data themselves are quite informative. Through an extensive forecasting experiment (for a range of daily and weekly data on stock markets and exchange rates) in chapter 2, we also demonstrate that linear time series models simply do not yield reliable forecasts. Of course, this does not automatically imply that nonlinear time series models might, but it is worth a try. As there is a host of possible nonlinear time series models, we review only what we believe are currently the most relevant ones and the ones we think are most likely to persist as practical descriptive and forecasting devices.

## 1.1 Introduction and outline of the book

Forecasting future returns on assets such as stocks and currencies' exchange rates is of obvious interest in empirical finance. For example, if one were able to forecast tomorrow's return on the Dow Jones index with some degree of precision, one could use this information in an investment decision today. Of course, we are seldom able to generate a very accurate prediction for asset returns, but hopefully we can perhaps at least forecast, for example, the sign of tomorrow's return.

Cambridge University Press

0521770416 - Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses and Dick Van Dijk

Excerpt

[More information](#)

## 2 Nonlinear time series models in empirical finance

The trade-off between return and risk plays a prominent role in many financial theories and models, such as Modern Portfolio theory and option pricing. Given that volatility is often regarded as a measure of this risk, one is interested not only in obtaining accurate forecasts of returns on financial assets, but also in forecasts of the associated volatility. Much recent evidence shows that volatility of financial assets is not constant, but rather that relatively volatile periods alternate with more tranquil ones. Thus, there may be opportunities to obtain forecasts of this time-varying risk.

Many models that are commonly used in empirical finance to describe returns and volatility are linear. There are, however, several indications that nonlinear models may be more appropriate (see section 1.2 for details). In this book, we therefore focus on the construction of nonlinear time series models that can be useful for describing and forecasting returns and volatility. While doing this, we do not aim to treat those models as ‘black boxes’. On the contrary, we provide ample details of representation and inference issues. Naturally, we will compare the descriptive models and their implied forecasts with those of linear models, in order to illustrate their potential relevance.

We focus on forecasting out-of-sample returns and volatility as such and abstain from incorporating such forecasts in investment strategies. We usually take (functions of) past returns as explanatory variables for current returns and volatility. With some degree of market efficiency, one may expect that most information is included in recent returns. Hence, we do not consider the possibility of explaining returns by variables that measure aspects of the underlying assets – such as, for example, specific news events and key indicators of economic activity. Another reason for restricting the analysis to univariate models is that we focus mainly on short-term forecasting – that is, not more than a few days or weeks ahead. Explanatory variables such as dividend yields, term structure variables and macroeconomic variables have been found mainly useful for predicting stock returns at longer horizons, ranging from one quarter to several years (see Kaul, 1996, for an overview of the relevant literature).

Numerous reasons may be evinced for the interest in nonlinear models. For example, in empirical finance it is by now well understood that financial time series data display *asymmetric behaviour*. An example of this behaviour is that large negative returns appear more frequently than large positive returns. Indeed, the stock market crash on Monday 19 October 1987 concerned a return of about  $-23$  per cent on the S&P 500 index, while for most stock markets we rarely observe positive returns of even 10 per cent or higher. Another example is that large negative returns are often a prelude to a period of substantial volatility, while large positive returns are less so. Needless to say, such asymmetries should be incorporated in a time series model used for description and out-of-sample forecasting, otherwise one may obtain forecasts that are always too low or too high. We will call such time series models,

Cambridge University Press

0521770416 - Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses and Dick Van Dijk

Excerpt

[More information](#)

## Introduction 3

which allow for an explicit description of asymmetries, *nonlinear time series models*.

An important debate in empirical finance concerns the question whether large negative returns, such as the 1987 stock market crash, are events that are atypical or naturally implied by an underlying process, which the nonlinear time series model should capture. It is well known that neglected *atypical events* can blur inference in linear time series models and can thus be the culprit of rather inaccurate forecasts. As nonlinear time series models are typically designed to accommodate features of the data that cannot be captured by linear models, one can expect that neglecting such atypical observations will have even more impact on out-of-sample forecasts. Therefore, in this book we pay quite considerable attention to take care of such observations while constructing nonlinear models.

Most descriptive and forecasting models in this book concern univariate financial time series – that is, we construct separate models for, for example, the Dow Jones and the FTSE index, ignoring the potential links between these two important stock markets. A multivariate model for the returns or volatilities of two or more stock markets jointly while allowing for asymmetries is a possible next step once univariate models have been considered. In specific sections in relevant chapters, we will give some attention to multivariate nonlinear models. It must be stressed, though, that the theory of multivariate nonlinear time series models has not yet been fully developed, and so we limit our discussion to only a few specific models.

This book is divided into six chapters. The current chapter and chapter 2 offer a first glance at some typical features of many financial time series and deal with some elementary concepts in time series analysis, respectively. Chapter 2 reviews only the key concepts needed for further reading, and the reader should consult textbooks on time series analysis, such as Hamilton (1994), Fuller (1996), Brockwell and Davis (1997) and Franses (1998), among others, for more detailed treatments. The concepts in chapter 2 can be viewed as the essential tools necessary for understanding the material in subsequent chapters. Readers who already are acquainted with most of the standard tools of time series analysis can skip this chapter and proceed directly to chapter 3.

Many economic time series display one or more of the following five features: a trend, seasonality, atypical observations, clusters of outliers and nonlinearity (see Franses, 1998). In this book, we focus on the last three features, while considering financial time series. The purpose of section 1.2 is to describe some of the characteristic features of financial time series, which strongly suggest the necessity for considering nonlinear time series models instead of linear models. In particular, we show that (1) large returns (in absolute terms) occur more frequently than one might expect under the assumption that the data are normally distributed (which often goes hand-in-hand with the use of linear

Cambridge University Press

0521770416 - Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses and Dick Van Dijk

Excerpt

[More information](#)**4 Nonlinear time series models in empirical finance**

models and which often is assumed in financial theory), (2) such large absolute returns tend to appear in clusters (indicating the possible presence of time-varying risk or volatility), (3) large negative returns appear more often than large positive ones in stock markets, while it may be the other way around for exchange rates, and (4) volatile periods are often preceded by large negative returns. The empirical analysis relies only on simple statistical techniques, and aims merely at highlighting which features of financial time series suggest the potential usefulness of, and should be incorporated in, a nonlinear time series model. For returns, features (1) and (3) suggest the usefulness of models that have different regimes (see also Granger, 1992). Those models will be analysed in detail in chapter 3 (and to some extent also in chapter 5). Features (2) and (4) suggest the relevance of models that allow for a description of time-varying volatility, with possibly different impact of positive and negative past returns. These models are the subject of chapter 4. A final feature of returns, which will be discussed at length in section 2.3, is that linear time series models do not appear to yield accurate out-of-sample forecasts, thus providing a more pragmatic argument for entertaining nonlinear models.

As running examples throughout this book, we consider daily indexes for eight major stock markets (including those of New York, Tokyo, London and Frankfurt), and eight daily exchange rates *vis-à-vis* the US dollar (including the Deutschmark and the British pound). We do not use all data to illustrate all models and methods, and often we take only a few series for selected applications. For convenience, we will analyse mainly the daily data in temporally aggregated form – that is, we mainly consider weekly data. In our experience, however, similar models can be useful for data sampled at other frequencies. As a courtesy to the reader who wishes to experiment with specific models, all data used in this book can be downloaded from (<http://www.few.eur.nl/few/people/franses>).

Chapter 3 focuses on nonlinear models for returns that impose a regime-switching structure. We review models with two or more regimes, models where the regimes switch abruptly and where they do not and models in which the switches between the different regimes are determined by specific functions of past returns or by an unobserved process. We pay attention to the impact of atypical events, and we show how these events can be incorporated in the model or in the estimation method, using a selective set of returns to illustrate the various models. In the last section of chapter 3 (3.7), we touch upon the issue of multivariate nonlinear models. The main conclusion from the empirical results in chapter 3 is that nonlinear models for returns may sometimes outperform linear models (in terms of within-sample fit and out-of-sample forecasting).

In chapter 4, we discuss models for volatility. We limit attention to those models that consider some form of autoregressive conditional heteroscedasticity (ARCH), although we briefly discuss the alternative class of stochastic volatility

Cambridge University Press

0521770416 - Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses and Dick Van Dijk

Excerpt

[More information](#)

## Introduction 5

models as well. The focus is on the basic ARCH model (which itself can be viewed as a nonlinear time series model) as was proposed in Engle (1982), and on testing, estimation, forecasting and the persistence of shocks. Again, we pay substantial attention to the impact of atypical events on estimated volatility. We also discuss extensions of the class of ARCH models in order to capture the asymmetries described in section 1.2. Generally, such extensions amount to modifying the standard ARCH model to allow for regime-switching effects in the persistence of past returns on future volatility.

Chapter 5 deals with models that allow the data to determine if there are different regimes that need different descriptive measures, while the number of regimes is also indicated by the data themselves. These flexible models are called ‘artificial neural network models’. In contrast to the prevalent strategy in the empirical finance literature (which may lead people to believe that these models are merely a passing fad), we decide, so to say, to ‘open up the black box’ and to explicitly demonstrate how and why these models can be useful in practice. Indeed, the empirical applications in this chapter suggest that neural networks can be quite useful for out-of-sample forecasting and for recognizing a variety of patterns in the data. We discuss estimation and model selection issues, and we pay attention to how such neural networks handle atypical observations.

Finally, chapter 6 contains a brief summary and some thoughts and suggestions for further research.

All computations in this book have been performed using GAUSS, version 3.2.35. The code of many of the programs that have been used can be downloaded from (<http://www.few.eur.nl/few/people/franses>).

In the remainder of this chapter we will turn our focus to some typical features of financial time series which suggest the potential relevance of nonlinear time series models.

## 1.2 Typical features of financial time series

Empirical research has brought forth a considerable number of stylized facts of high-frequency financial time series. The purpose of this section is to describe some of these characteristic features. In particular, we show that returns on financial assets display erratic behaviour, in the sense that large outlying observations occur with rather high-frequency, that large negative returns occur more often than large positive ones, that these large returns tend to occur in clusters and that periods of high volatility are often preceded by large negative returns. Using simple and easy-to-compute statistical and graphical techniques, we illustrate these properties for a number of stock index and exchange rate returns, sampled at daily and weekly frequencies. The data are described in more detail below. Throughout this section we emphasize that the above-mentioned stylized facts seem to imply the necessity of considering nonlinear models to describe

6 Nonlinear time series models in empirical finance

the observed patterns in such financial time series adequately and to render sensible out-of-sample forecasts. In chapter 2, we will show more rigorously that linear models appear not to be useful for out-of-sample forecasting of returns on financial assets.

Finally, it should be remarked that the maintained hypothesis for high-frequency financial time series is that (logarithmic) prices of financial assets display random walk-type behaviour (see Campbell, Lo and MacKinlay, 1997). Put differently, when linear models are used, asset prices are assumed to conform to a martingale – that is, the expected value of (the logarithm of) tomorrow’s price  $P_{t+1}$ , given all relevant information up to and including today, denoted as  $\Omega_t$ , should equal today’s value, possibly up to a deterministic growth component which is denoted as  $\mu$ , or,

$$E[\ln P_{t+1} | \Omega_t] = \ln P_t + \mu, \tag{1.1}$$

where  $E[\cdot]$  denotes the mathematical expectation operator and  $\ln$  denotes the natural logarithmic transformation. In section 2.3 we will examine if (1.1) also gives the best forecasts when compared with other linear models.

*The data*

The data that we use to illustrate the typical features of financial time series consist of eight indexes of major stock markets and eight exchange rates *vis-à-vis* the US dollar. To be more precise, we employ the indexes of the stock markets in Amsterdam (EOE), Frankfurt (DAX), Hong Kong (Hang Seng), London (FTSE100), New York, (S&P 500), Paris (CAC40), Singapore (Singapore All Shares) and Tokyo (Nikkei). The exchange rates are the Australian dollar, British pound, Canadian dollar, German Deutschmark, Dutch guilder, French franc, Japanese yen and Swiss franc, all expressed as a number of units of the foreign currency per US dollar. The sample period for the stock indexes runs from 6 January 1986 until 31 December 1997, whereas for the exchange rates the sample covers the period from 2 January 1980 until 31 December 1997. The original series are sampled at daily frequency. The sample periods correspond with 3,127 and 4,521 observations for the stock market indexes and exchange rates, respectively. We often analyse the series on a weekly basis, in which case we use observations recorded on Wednesdays. The stock market data have been obtained from Datastream, whereas the exchange rate data have been obtained from the New York Federal Reserve.

Figures 1.1 and 1.2 offer a first look at the data by showing a selection of the original price series  $P_t$  and the corresponding logarithmic returns measured in percentage terms, denoted  $y_t$  and computed as

$$y_t = 100 \cdot (p_t - p_{t-1}), \tag{1.2}$$

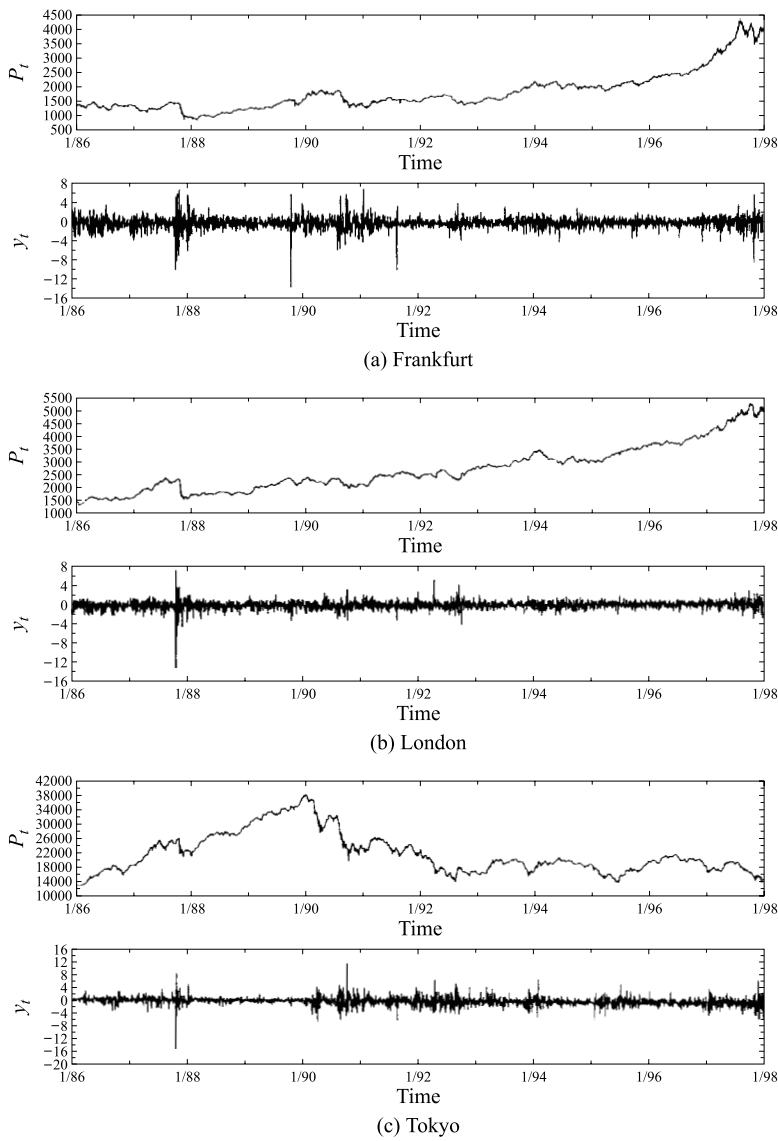


Figure 1.1 Daily observations on the level (upper panel) and returns (lower panel) of (a) the Frankfurt, (b) the London and (c) the Tokyo stock indexes, from 6 January 1986 until 31 December 1997

8 Nonlinear time series models in empirical finance

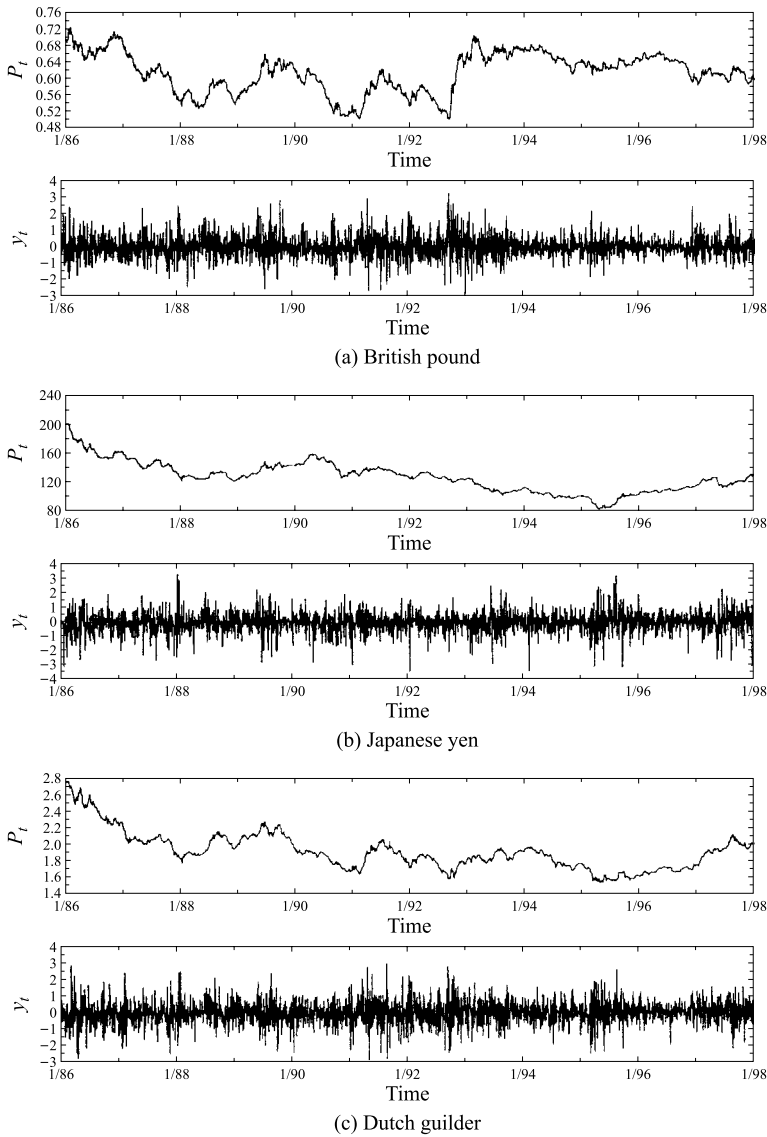


Figure 1.2 Daily observations on the level (upper panel) and returns (lower panel) of (a) the British pound, (b) the Japanese yen and (c) the Dutch guilder exchange rates *vis-à-vis* the US dollar, from 6 January 1986 until 31 December 1997



where  $p_t = \ln(P_t)$ . Strictly speaking, returns should also take into account dividends, but for daily data one often uses (1.2). Prices and returns for the Frankfurt, London and Tokyo indexes are shown in figure 1.1, and prices and returns for the British pound, Japanese yen and Dutch guilder exchange rates are shown in figure 1.2 (also for the period 1986–97).

Summary statistics for the stock and exchange rate returns are given in tables 1.1 and 1.2, respectively, for both daily and weekly sampling frequencies. These statistics are used in the discussion of the characteristic features of these series below.

*Large returns occur more often than expected*

One of the usual assumptions in the (theoretical) finance literature is that the logarithmic returns  $y_t$  are normally distributed random variables, with mean  $\mu$  and variance  $\sigma^2$ , that is,

$$y_t \sim N(\mu, \sigma^2). \tag{1.3}$$

Table 1.1 *Summary statistics for stock returns*

Stock market	Mean	Med	Min	Max	Var	Skew	Kurt
<i>Daily returns</i>							
Amsterdam	0.038	0.029	−12.788	11.179	1.279	−0.693	19.795
Frankfurt	0.035	0.026	−13.710	7.288	1.520	−0.946	15.066
Hong Kong	0.057	0.022	−40.542	17.247	2.867	−5.003	119.241
London	0.041	0.027	−13.029	7.597	0.845	−1.590	27.408
New York	0.049	0.038	−22.833	8.709	0.987	−4.299	99.680
Paris	0.026	0.000	−10.138	8.225	1.437	−0.529	10.560
Singapore	0.019	0.000	−9.403	14.313	1.021	−0.247	28.146
Tokyo	0.005	0.000	−16.135	12.430	1.842	−0.213	14.798
<i>Weekly returns</i>							
Amsterdam	0.190	0.339	−19.962	7.953	5.853	−1.389	11.929
Frankfurt	0.169	0.354	−18.881	8.250	6.989	−1.060	8.093
Hong Kong	0.283	0.556	−34.969	11.046	13.681	−2.190	18.258
London	0.207	0.305	−17.817	9.822	4.617	−1.478	15.548
New York	0.246	0.400	−16.663	6.505	4.251	−1.370	11.257
Paris	0.128	0.272	−20.941	11.594	8.092	−0.995	9.167
Singapore	0.091	0.110	−27.335	10.510	6.986	−2.168	23.509
Tokyo	0.025	0.261	−10.892	12.139	8.305	−0.398	4.897

Notes: Summary statistics for returns on stock market indexes.  
The sample period is 6 January 1986 until 31 December 1997, which equals 3,127 (625) daily (weekly) observations.

10 Nonlinear time series models in empirical finance

Table 1.2 Summary statistics for exchange rate returns

Currency	Mean	Med	Min	Max	Var	Skew	Kurt
<i>Daily returns</i>							
Australian dollar	0.012	−0.012	−5.074	10.554	0.377	1.893	35.076
British pound	0.006	0.000	−4.589	3.843	0.442	0.058	5.932
Canadian dollar	0.006	0.000	−1.864	1.728	0.076	0.101	6.578
Dutch guilder	−0.000	0.012	−3.985	3.188	0.464	−0.143	4.971
French franc	0.008	0.016	−3.876	5.875	0.457	0.054	6.638
German Dmark	−0.001	0.017	−4.141	3.227	0.475	−0.136	4.921
Japanese yen	−0.016	0.006	−5.630	3.366	0.478	−0.541	6.898
Swiss franc	−0.003	0.020	−4.408	3.300	0.582	−0.188	4.557
<i>Weekly returns</i>							
Australian dollar	0.057	−0.022	−5.526	10.815	1.731	1.454	11.906
British pound	0.033	−0.027	−7.397	8.669	2.385	0.218	6.069
Canadian dollar	0.022	0.016	−2.551	2.300	0.343	0.040	4.093
Dutch guildier	0.007	0.051	−7.673	7.212	2.416	−0.155	4.518
French franc	0.043	0.074	−7.741	6.858	2.383	−0.014	5.006
German Dmark	0.005	0.052	−8.113	7.274	2.483	−0.168	4.545
Japanese yen	−0.064	0.059	−6.546	6.582	2.192	−0.419	4.595
Swiss franc	−0.008	0.105	−7.969	6.636	2.929	−0.314	3.930

Notes: Summary statistics for exchange rate returns.  
The sample period is 2 January 1980 until 31 December 1997, which equals 4,521 (939) daily (weekly) observations.

The *kurtosis* of  $y_t$  is defined as

$$K_y = E \left[ \frac{(y_t - \mu)^4}{\sigma^4} \right]. \tag{1.4}$$

For an observed time series  $y_1, \dots, y_n$ , the kurtosis can be estimated consistently by the sample analogue of (1.4),

$$\widehat{K}_y = \frac{1}{n} \sum_{t=1}^n \frac{(y_t - \hat{\mu})^4}{\hat{\sigma}^4}, \tag{1.5}$$

where  $\hat{\mu} = \frac{1}{n} \sum_{t=1}^n y_t$  and  $\hat{\sigma}^2 = \frac{1}{n} \sum_{t=1}^n (y_t - \hat{\mu})^2$  are the sample mean and variance, respectively. The kurtosis for the normal distribution is equal to 3. One of the features which stands out most prominently from the last columns of tables 1.1 and 1.2 is that the kurtosis of all series is much larger than this normal value, especially for the daily series. This reflects the fact that the tails of the