

Nonlinear Time Series Models in Empirical Finance

Philip Hans Franses

and

Dick van Dijk



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK

www.cup.cam.ac.uk

40 West 20th Street, New York, NY 10011-43211, USA

www.cup.org

10 Stamford Road, Oakleigh, Melbourne 3166, Australia

Ruiz de Alarcón 13, 28014 Madrid, Spain

© Franses and van Dijk 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United Kingdom at the University Press, Cambridge

Typeset Times 10/12pt. System LaTeX

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Franses, Philip Hans, 1963–

Nonlinear time series models in empirical finance/Philip Hans Franses, Dick van Dijk.
p. cm.

Includes bibliographical references and index.

ISBN 0 521 77041 6 ISBN 0 521 77965 0 (pbk.)

1. Finance—Mathematical models. 2. Time-series analysis. I. Dijk, Dick van. II. Title.

HG106 .F73 2000 332'.01'5118—dc21 99-088504

ISBN 0521 77041 6 hardback

ISBN 0521 77965 0 paperback

Contents

<i>List of figures</i>	<i>page</i> ix
<i>List of tables</i>	xi
<i>Preface</i>	xv
1 Introduction	1
1.1 Introduction and outline of the book	1
1.2 Typical features of financial time series	5
2 Some concepts in time series analysis	20
2.1 Preliminaries	20
2.2 Empirical specification strategy	27
2.3 Forecasting returns with linear models	44
2.4 Unit roots and seasonality	51
2.5 Aberrant observations	61
3 Regime-switching models for returns	69
3.1 Representation	71
3.2 Estimation	83
3.3 Testing for regime-switching nonlinearity	100
3.4 Diagnostic checking	108
3.5 Forecasting	117
3.6 Impulse response functions	125
3.7 On multivariate regime-switching models	132
4 Regime-switching models for volatility	135
4.1 Representation	136
4.2 Testing for GARCH	157
4.3 Estimation	170

4.4	Diagnostic checking	182
4.5	Forecasting	187
4.6	Impulse response functions	197
4.7	On multivariate GARCH models	200
5	Artificial neural networks for returns	206
5.1	Representation	207
5.2	Estimation	215
5.3	Model evaluation and model selection	222
5.4	Forecasting	234
5.5	ANNs and other regime-switching models	237
5.6	Testing for nonlinearity using ANNs	245
6	Conclusions	251
	<i>Bibliography</i>	254
	<i>Author index</i>	272
	<i>Subject index</i>	277

Figures

1.1	Stock indexes – levels and returns	<i>page</i> 7
1.2	Exchange rates – levels and returns	8
1.3	Distributions of stock index returns	11
1.4	Distributions of exchange rate returns	12
1.5	Scatterplot of daily returns on the Amsterdam stock index	14
1.6	Scatterplot of daily returns on the Frankfurt stock index	15
1.7	Scatterplot of daily returns on the London stock index	16
1.8	Scatterplot of daily returns on the British pound	17
1.9	Scatterplot of daily returns on the Canadian dollar	18
1.10	Scatterplot of daily returns on the Dutch guilder	19
2.1	Autocorrelations of stock index returns	31
2.2	Autocorrelations of exchange rate returns	32
2.3	Additive and innovative outliers in an AR(1) model	63
2.4	Weight functions for robust estimation	67
3.1	Logistic functions	72
3.2	Realizations from a SETAR model	73
3.3	Scatterplots for realizations from a SETAR model	74
3.4	Sequences of LR-statistics for realizations from a SETAR model	86
3.5	Absolute weekly returns on the Frankfurt stock index and regime probabilities in a Markov-Switching model	97
3.6	Weekly returns on the Dutch guilder exchange rate and weights from robust estimation of a SETAR model	99
3.7	Transition function in a STAR model for returns on the Dutch guilder exchange rate	109
3.8	Transition function in a STAR model for absolute returns on the Tokyo stock index	111
3.9	Conditional distributions for a SETAR model	123
3.10	Generalized impulse responses in a STAR model for returns on the Dutch guilder exchange rate	131

3.11	Generalized impulse responses in a STAR model for returns on the Dutch guilder exchange rate	131
4.1	News impact curves for nonlinear GARCH models	150
4.2	News impact curve for the ANST-GARCH model	154
4.3	Conditional standard deviation in nonlinear GARCH models for returns on the Tokyo stock index	176
4.4	Conditional standard deviation in GARCH(1,1) models for weekly stock index and exchange rate returns	179
5.1	Skeleton of an ANN	209
5.2	Structure of the hidden layer in an ANN	211
5.3	Architecture of the single hidden layer feedforward ANN	213
5.4	ANN and additive outliers	215
5.5	ANN and innovative outliers	216
5.6	ANN and level shifts	217
5.7	Output of hidden units in an ANN for returns on the Japanese yen exchange rate	224
5.8	Skeleton of an ANN for returns on the Japanese yen exchange rate	225
5.9	Output of hidden units in an ANN for absolute returns on the Frankfurt stock index	227
5.10	Skeleton of an ANN for absolute returns on the Frankfurt stock index	228
5.11	Moving averages for returns on the Japanese yen exchange rate	229
5.12	Output–input derivatives in an ANN for returns on the Japanese yen exchange rate	233
5.13	Output–input derivatives in an ANN for absolute returns on the Frankfurt stock index	233
5.14	Impulse responses in an ANN for absolute returns on the Frankfurt stock index	238

Tables

1.1	Summary statistics for stock returns	<i>page</i> 9
1.2	Summary statistics for exchange rate returns	10
1.3	Correlation between squared returns at day t and returns at day $t - 1$	18
2.1	Average ranks of linear models to forecast stock returns according to MSPE, 1991–1997	46
2.2	Average ranks of linear models to forecast stock returns according to MAPE, 1991–1997	47
2.3	Average ranks of linear models to forecast stock returns according to MedSPE, 1991–1997	48
2.4	Forecast comparison of linear models with random walk – stock returns, squared prediction errors, 1991–1997	49
2.5	Forecast comparison of linear models with random walk – stock returns, absolute prediction errors, 1991–1997	50
2.6	Performance of linear models in forecasting sign of stock returns, 1991–1997	52
2.7	Daily means and variances of stock index returns	59
2.8	Periodic autocorrelations of stock returns	61
3.1	AIC for $AR(p)$ models estimated on simulated SETAR series	78
3.2	AIC values for SETAR models for weekly returns on the Dutch guilder exchange rate	88
3.3	SETAR estimates for weekly percentage returns on the Dutch guilder exchange rate	89
3.4	Parameter estimates for a MSW model for weekly absolute returns on the Frankfurt stock index	96
3.5	p -values for HCC test of linearity against a SETAR alternative for weekly returns on the Dutch guilder exchange rate	106
3.6	p -values of LM-type test for STAR nonlinearity for weekly returns on the Dutch guilder exchange rate	107

3.7	Parameter estimates for a STAR model for weekly returns on the Dutch guilder exchange rate	108
3.8	p -values of LM-type test for STAR nonlinearity for weekly absolute returns on the Tokyo stock index	110
3.9	Parameter estimates for a STAR model for weekly absolute returns on the Tokyo stock index	110
3.10	Diagnostic tests of a STAR model estimated for weekly returns on the Dutch guilder exchange rate	115
3.11	Diagnostic tests of a STAR model estimated for absolute weekly returns on the Tokyo stock index	116
3.12	Forecast evaluation of a STAR model for weekly returns on the Dutch guilder exchange rate	126
3.13	Forecast evaluation of a STAR model for weekly returns on the Dutch guilder exchange rate, 1990–1997	127
4.1	Testing for ARCH in weekly stock index returns	158
4.2	Testing for ARCH in weekly exchange rate returns	159
4.3	Testing for asymmetric ARCH effects in weekly stock index and exchange rate returns	161
4.4	Testing for nonlinear ARCH in weekly stock index returns	163
4.5	Testing for nonlinear ARCH in weekly exchange rate returns	164
4.6	Testing for ARCH and QARCH in simulated SETAR series	165
4.7	Rejection frequencies of standard and robust tests for (nonlinear) ARCH in the presence of outliers	167
4.8	Properties of standard and robust tests for ARCH in the presence of patchy outliers	169
4.9	Estimates of nonlinear GARCH(1,1) models for weekly returns on the Tokyo stock index	175
4.10	Estimates of GARCH(1,1) models for weekly stock index and exchange rate returns	177
4.11	Percentiles of the distribution of the outlier detection statistic in GARCH(1,1) models	182
4.12	Estimates of GARCH(1,1) models for weekly returns on the Amsterdam and New York stock indexes, before and after outlier correction	183
4.13	Diagnostic tests for estimated GARCH models for weekly stock index and exchange rate returns	187
4.14	Forecast evaluation of nonlinear GARCH models for weekly returns on the Tokyo stock index, as compared to the GARCH (1,1) model	197
4.15	Forecast evaluation of nonlinear GARCH models for weekly returns on the Tokyo stock index	198

4.16	Testing for common ARCH effects in weekly stock index and exchange rate returns	205
5.1	Performance of ANNs when series are generated from an AR(2) model contaminated with AOs	218
5.2	Performance of ANNs when series are generated from an AR(2) model contaminated with IOs	219
5.3	Performance of ANNs applied to weekly returns on the Japanese yen exchange rate	223
5.4	Performance of ANNs applied to absolute weekly returns on the Frankfurt stock index	226
5.5	Performance of ANNs with technical trading rule applied to weekly returns on Japanese yen	231
5.6	Performance of ANNs when series are generated from a SETAR model	242
5.7	Performance of ANNs when series are generated from a Markov-Switching model	243
5.8	Performance of ANNs when series are generated from a bilinear model	244
5.9	Performance of ANNs when series are generated from a GARCH(1,1) model	245
5.10	Testing for nonlinearity in weekly stock index and exchange rate returns with ANN-based tests	248
5.11	Testing for nonlinearity in weekly absolute stock index and exchange rate returns with ANN-based tests	249

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.

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,

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

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

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

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 Ω_t , should equal today's value, possibly up to a deterministic growth component which is denoted as μ , or,

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

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

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




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




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 μ and variance σ^2 , that is,

$$y_t \sim N(\mu, \sigma^2). \quad (1.3)$$

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.

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 guilder	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

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

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

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

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

distributions of these series are fatter than the tails of the normal distribution. Put differently, large observations occur (much) more often than one might expect for a normally distributed variable.

This is illustrated further in figures 1.3 and 1.4, which show estimates of the distributions $f(y)$ of the daily returns on the Frankfurt and London stock indexes




Figure 1.3 Kernel estimates of the distribution of daily returns on (a) the Frankfurt and (b) the London stock indexes (solid line) and normal distribution with same mean and variance (dashed line); each whisker represents one observation




Figure 1.4 Kernel estimates of the distribution of daily returns on (a) the British pound and (b) Japanese yen exchange rates *vis-à-vis* the US dollar (solid line) and normal distribution with same mean and variance (dashed line); each whisker represents one observation

and the British pound and Japanese yen exchange rates, respectively. The estimates are obtained with a kernel density estimator,

$$\hat{f}(y) = \frac{1}{nh} \sum_{t=1}^n K\left(\frac{y_t - y}{h}\right),$$

where $K(z)$ is a function which satisfies $\int K(z) dz = 1$ and h is the so-called bandwidth. Usually $K(z)$ is taken to be a unimodal probability density function; here we use the Gaussian kernel

$$K(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}z^2\right).$$

Following Silverman (1986), we set the bandwidth h according to $h = 0.9 \cdot \min(\hat{\sigma}, \text{iqr}/1.349)n^{-1/5}$, where iqr denotes the sample interquartile range – that is, $\text{iqr} = y_{\lfloor (3n/4) \rfloor} - y_{\lfloor (n/4) \rfloor}$, where $y_{(i)}$ is the i th order statistic of the series y_t , $t = 1, \dots, n$, and $\lfloor \cdot \rfloor$ denotes the integer part. (See Wand and Jones, 1995, for discussion of this and other kernel estimators, and various methods of bandwidth selection.) In all graphs, a normal distribution with mean and variance obtained from tables 1.1 and 1.2 for the different series has also been drawn for ease of comparison. Each whisker on the horizontal axis represents one observation. Clearly, all distributions are more peaked and have fatter tails than the corresponding normal distributions. Thus, both very small and very large observations occur more often compared to a normally distributed variable with the same first and second moments.

Finally, it is worth noting that the kurtosis of the stock returns is much larger than the kurtosis of the exchange rate returns, at both the daily and weekly sampling frequency. This may reflect the fact that central banks can intervene in the foreign exchange market, while there are virtually no such opportunities in stock markets.

Large stock market returns are often negative

The skewness of y_t is defined as

$$\text{SK}_y = E \left[\frac{(y_t - \mu)^3}{\sigma^3} \right], \quad (1.6)$$

and is a measure of the asymmetry of the distribution of y_t . The skewness for an observed time series $y_1, \dots, y_n$ can be estimated consistently by the sample analogue of (1.6) as

$$\widehat{\text{SK}}_y = \frac{1}{n} \sum_{t=1}^n \frac{(y_t - \hat{\mu})^3}{\hat{\sigma}^3}. \quad (1.7)$$

All symmetric distributions, including the normal distribution, have skewness equal to zero. From table 1.1 it is seen that the stock return series all have negative skewness, which implies that the left tail of the distribution is fatter than the right tail, or that large negative returns tend to occur more often than

large positive ones. This is visible in the distributions displayed in figure 1.3 as well, as more whiskers are present in the left tail than in the right tail.

Skewness of the daily exchange rate returns is positive for certain currencies, while it is negative for others. This makes sense, as it is not *a priori* clear why exchange rate returns should have either positive or negative skewness when measured in the way we do here.

Large returns tend to occur in clusters

From figures 1.1 and 1.2 it appears that relatively volatile periods, characterized by large price changes – and, hence, large returns – alternate with more tranquil periods in which prices remain more or less stable and returns are, consequently, small. In other words, large returns seem to occur in clusters. This feature of our time series becomes even more apparent when inspecting scatterplots of the return of day t , denoted y_t , against the return of day $t - 1$. Figures 1.5–1.7 provide such plots for the daily observed Amsterdam, Frankfurt and London stock indexes. Similar scatterplots for daily data on the British




Figure 1.5 Scatterplot of the return on the Amsterdam stock index on day t , y_t , against the return on day $t - 1$

The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses




Figure 1.6 Scatterplot of the return on the Frankfurt stock index on day t , y_t , against the return on day $t - 1$

The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses

pound, Canadian dollar and Dutch guilder are shown in figures 1.8–1.10. In these scatterplots, the observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves. All observations that are starting and/or end-points of arrows are marked with crosses.

Following the route indicated by the arrows reveals that the return series frequently travel around the main cloud of observations for an extended period of time. This holds for stock returns in particular. For example, the arrows in figure 1.5 really comprise only two stretches of large returns. The first stretch starts at $(y_{t-1}, y_t) = (-1.90, -2.15)$ which corresponds to 14 and 15 October 1987. On subsequent trading days, the return on the Amsterdam stock index was equal to $-0.58, -12.79$ (19 October), $-6.10, 8.81, -7.52, 0.22, -9.74, 3.21, -6.14$ and 0.25 per cent on 29 October, which is where the first path ends. The second one starts with the pair of returns on 6 and 9 November 1987, which are




Figure 1.7 Scatterplot of the return on the London stock index on day t , y_t , against the return on day $t - 1$

The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses

equal to 0.31 and -7.39 , respectively, and were followed by returns of -4.52 , 11.18 , 8.35 , -2.79 , 3.89 and -3.99 per cent.

From figures 1.8–1.10 it appears that clustering of large returns occurs less frequently for exchange rates. The arrows seem to constitute a three-cycle quite often, where ‘three-cycle’ refers to the situation where the return series leaves the main cloud of observations owing to a large value of y_t , moves to the next observation (which necessarily is outside of the main cloud as well, as now y_{t-1} is large) and moves back into the main clutter the next day. Evidently, such three-cycles are caused by a single large return. Still, some longer stretches of arrows are present as well.

Large volatility often follows large negative stock market returns

Another property of the stock return series that can be inferred from the scatterplots presented is that periods of large volatility tend to be triggered by a




Figure 1.8 Scatterplot of the return on the British pound/US dollar exchange rate on day t , y_t , against the return on day $t - 1$

The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses

large negative return. Further inspection of figures 1.5–1.7 shows that the stock return series almost invariably leave the central cloud in a southern direction – that is, today’s return is large and negative. Given that it can take quite some time before the return series calms down and that scatter observations disappear into the main cloud again, it seems justified to state that a volatile period often starts with a large negative return.

The second column of table 1.3 contains estimates of the correlation between the squared return at day t and the return at day $t - 1$ for the various stock indexes. The fact that all these correlations are negative also illustrates that large volatility often follows upon a negative return.

For the exchange rate returns this property is much less clear-cut (as it should be, as the return series can be inverted by simply expressing the exchange rate as the number of US dollars per unit of foreign currency). Figures 1.8–1.10 do not reveal any preference of the exchange rate return series to leave the main cloud of observations either to the north or to the south. The estimates of the

Table 1.3 *Correlation between squared returns at day t and returns at day $t - 1$*

Stock market	$\text{Corr}(y_t^2, y_{t-1})$	Exchange rate	$\text{Corr}(y_t^2, y_{t-1})$
Amsterdam	-0.049	Australian dollar	0.168
Frankfurt	-0.095	British pound	0.074
Hong Kong	-0.081	Canadian dollar	0.041
London	-0.199	Dutch guilder	0.042
New York	-0.108	French franc	0.047
Paris	-0.042	German Dmark	0.041
Singapore	-0.107	Japanese yen	-0.008
Tokyo	-0.130	Swiss franc	0.014

Note: Correlation between squared return at day t and return at day $t - 1$ for stock market indices and exchange rates.




Figure 1.9 Scatterplot of the return on the Canadian dollar/US dollar exchange rate on day t , y_t , against the return on day $t - 1$
The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses




Figure 1.10 Scatterplot of the return on the Dutch guilder exchange rate on day t , y_t , against the return on day $t - 1$

The observations for the three largest negative and the three largest positive values of y_t are connected with the two preceding and the two following observations by means of arrows, pointing in the direction in which the time series evolves; all observations that are starting- and/or end-points of arrows are marked with crosses

correlations between y_t^2 and y_{t-1} , as shown in the final column of table 1.3, are positive for all exchange rate series except the Japanese yen.

To summarize, the typical features of financial time series documented in this first chapter seem to require nonlinear models, simply because linear models would not be able to generate data that have these features. Before we turn to a discussion of nonlinear models for the returns in chapter 3, we first review several important time series analysis tools in chapter 2, which are needed for a better understanding of the material later in the book.