

CHRIS BROOKS

INTRODUCTORY
ECONOMETRICS
FOR FINANCE

3rd Edition

Introductory Econometrics for Finance

This bestselling and thoroughly classroom-tested textbook is a complete resource for finance students. A comprehensive and illustrated discussion of the most common empirical approaches in finance prepares students for using econometrics in practice, while detailed financial case studies help them understand how the techniques are used in relevant financial contexts. Worked examples from the latest version of the popular statistical software EViews guide students to implement their own models and interpret results. Learning outcomes, key concepts and end-of-chapter review questions (with full solutions online) highlight the main chapter takeaways and allow students to self-assess their understanding. Building on the successful data- and problem-driven approach of previous editions, this third edition has been updated with new data, extensive examples and additional introductory material on mathematics, making the book more accessible to students encountering econometrics for the first time. A companion website, with numerous student and instructor resources, completes the learning package.

Chris Brooks is Professor of Finance and Director of Research at the ICMA Centre, Henley Business School, University of Reading, UK where he also obtained his PhD. He has diverse research interests and has published over a hundred articles in leading academic and practitioner journals, and six books. He is Associate Editor of several journals, including the Journal of Business Finance and Accounting, the International Journal of Forecasting and the British Accounting Review. He acts as consultant and advisor for various banks, corporations and professional bodies in the fields of finance, real estate, and econometrics.

Introductory Econometrics for Finance

THIRD EDITION

Chris Brooks

The ICMA Centre, Henley Business School, University of Reading



CAMBRIDGE UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Published in the United States of America by Cambridge University Press, New York

Cambridge University Press is a part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781107661455

© Chris Brooks 2014

This publication 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 2002

Second edition 2008

Third edition published 2014

Printed in the United Kingdom by MPG Printgroup Ltd, Cambridge

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

Library of Congress Cataloguing in Publication data

ISBN 978-1-107-03466-2 Hardback

ISBN 978-1-107-66145-5 Paperback

Additional resources for this publication at www.cambridge.org/brooks

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.



Contents

List of figures	<i>page</i> xii
List of tables	xv
List of boxes	xvii
List of screenshots	xix
Preface to the third edition	xxi
Acknowledgements	xxv
1 Introduction	1
1.1 What is econometrics?	2
1.2 Is financial econometrics different from ‘economic econometrics’?	2
1.3 Types of data	4
1.4 Returns in financial modelling	7
1.5 Steps involved in formulating an econometric model	11
1.6 Points to consider when reading articles in empirical finance	12
1.7 A note on Bayesian versus classical statistics	13
1.8 An introduction to EViews	14
1.9 Further reading	24
1.10 Outline of the remainder of this book	24
2 Mathematical and statistical foundations	28
2.1 Functions	28
2.2 Differential calculus	37
2.3 Matrices	41
2.4 Probability and probability distributions	56
2.5 Descriptive statistics	61
3 A brief overview of the classical linear regression model	75
3.1 What is a regression model?	75
3.2 Regression versus correlation	76
3.3 Simple regression	76
3.4 Some further terminology	84
3.5 Simple linear regression in EViews – estimation of an optimal hedge ratio	86

3.6	The assumptions underlying the classical linear regression model	90
3.7	Properties of the OLS estimator	91
3.8	Precision and standard errors	93
3.9	An introduction to statistical inference	98
3.10	A special type of hypothesis test: the t -ratio	111
3.11	An example of a simple t -test of a theory in finance: can US mutual funds beat the market?	113
3.12	Can UK unit trust managers beat the market?	115
3.13	The overreaction hypothesis and the UK stock market	116
3.14	The exact significance level	120
3.15	Hypothesis testing in EViews – example 1: hedging revisited	121
3.16	Hypothesis testing in EViews – example 2: the CAPM	123
	Appendix: Mathematical derivations of CLRM results	127
4	Further development and analysis of the classical linear regression model	134
4.1	Generalising the simple model to multiple linear regression	134
4.2	The constant term	135
4.3	How are the parameters (the elements of the β vector) calculated in the generalised case?	137
4.4	Testing multiple hypotheses: the F -test	139
4.5	Sample EViews output for multiple hypothesis tests	144
4.6	Multiple regression in EViews using an APT-style model	145
4.7	Data mining and the true size of the test	150
4.8	Goodness of fit statistics	151
4.9	Hedonic pricing models	156
4.10	Tests of non-nested hypotheses	159
4.11	Quantile regression	161
	Appendix 4.1: Mathematical derivations of CLRM results	168
	Appendix 4.2: A brief introduction to factor models and principal components analysis	170
5	Classical linear regression model assumptions and diagnostic tests	179
5.1	Introduction	179
5.2	Statistical distributions for diagnostic tests	180
5.3	Assumption 1: $E(u_i) = 0$	181
5.4	Assumption 2: $\text{var}(u_i) = \sigma^2 < \infty$	181
5.5	Assumption 3: $\text{cov}(u_i, u_j) = 0$ for $i \neq j$	188
5.6	Assumption 4: the x_t are non-stochastic	208
5.7	Assumption 5: the disturbances are normally distributed	209
5.8	Multicollinearity	217
5.9	Adopting the wrong functional form	220
5.10	Omission of an important variable	224
5.11	Inclusion of an irrelevant variable	225

5.12	Parameter stability tests	226
5.13	Measurement errors	235
5.14	A strategy for constructing econometric models and a discussion of model-building philosophies	238
5.15	Determinants of sovereign credit ratings	240
6	Univariate time series modelling and forecasting	251
6.1	Introduction	251
6.2	Some notation and concepts	252
6.3	Moving average processes	256
6.4	Autoregressive processes	259
6.5	The partial autocorrelation function	266
6.6	ARMA processes	268
6.7	Building ARMA models: the Box–Jenkins approach	273
6.8	Constructing ARMA models in EViews	276
6.9	Examples of time series modelling in finance	281
6.10	Exponential smoothing	283
6.11	Forecasting in econometrics	285
6.12	Forecasting using ARMA models in EViews	296
6.13	Exponential smoothing models in EViews	298
7	Multivariate models	305
7.1	Motivations	305
7.2	Simultaneous equations bias	307
7.3	So how can simultaneous equations models be validly estimated?	308
7.4	Can the original coefficients be retrieved from the π s?	309
7.5	Simultaneous equations in finance	311
7.6	A definition of exogeneity	312
7.7	Triangular systems	314
7.8	Estimation procedures for simultaneous equations systems	315
7.9	An application of a simultaneous equations approach to modelling bid–ask spreads and trading activity	318
7.10	Simultaneous equations modelling using EViews	323
7.11	Vector autoregressive models	326
7.12	Does the VAR include contemporaneous terms?	332
7.13	Block significance and causality tests	333
7.14	VARs with exogenous variables	335
7.15	Impulse responses and variance decompositions	336
7.16	VAR model example: the interaction between property returns and the macroeconomy	338
7.17	VAR estimation in EViews	344
8	Modelling long-run relationships in finance	353
8.1	Stationarity and unit root testing	353
8.2	Tests for unit roots in the presence of structural breaks	365

8.3	Testing for unit roots in EViews	369
8.4	Cointegration	373
8.5	Equilibrium correction or error correction models	375
8.6	Testing for cointegration in regression: a residuals-based approach	376
8.7	Methods of parameter estimation in cointegrated systems	377
8.8	Lead-lag and long-term relationships between spot and futures markets	380
8.9	Testing for and estimating cointegrating systems using the Johansen technique based on VARs	386
8.10	Purchasing power parity	390
8.11	Cointegration between international bond markets	391
8.12	Testing the expectations hypothesis of the term structure of interest rates	398
8.13	Testing for cointegration and modelling cointegrated systems using EViews	400
9	Modelling volatility and correlation	415
9.1	Motivations: an excursion into non-linearity land	415
9.2	Models for volatility	420
9.3	Historical volatility	420
9.4	Implied volatility models	421
9.5	Exponentially weighted moving average models	421
9.6	Autoregressive volatility models	422
9.7	Autoregressive conditionally heteroscedastic (ARCH) models	423
9.8	Generalised ARCH (GARCH) models	428
9.9	Estimation of ARCH/GARCH models	431
9.10	Extensions to the basic GARCH model	439
9.11	Asymmetric GARCH models	440
9.12	The GJR model	440
9.13	The EGARCH model	441
9.14	GJR and EGARCH in EViews	441
9.15	Tests for asymmetries in volatility	443
9.16	GARCH-in-mean	445
9.17	Uses of GARCH-type models including volatility forecasting	446
9.18	Testing non-linear restrictions or testing hypotheses about non-linear models	452
9.19	Volatility forecasting: some examples and results from the literature	454
9.20	Stochastic volatility models revisited	461
9.21	Forecasting covariances and correlations	463
9.22	Covariance modelling and forecasting in finance: some examples	464
9.23	Simple covariance models	466
9.24	Multivariate GARCH models	467
9.25	Direct correlation models	471

9.26	Extensions to the basic multivariate GARCH model	472
9.27	A multivariate GARCH model for the CAPM with time-varying covariances	474
9.28	Estimating a time-varying hedge ratio for FTSE stock index returns	475
9.29	Multivariate stochastic volatility models	478
9.30	Estimating multivariate GARCH models using EViews	480
	Appendix: Parameter estimation using maximum likelihood	484
10	Switching models	490
10.1	Motivations	490
10.2	Seasonalities in financial markets: introduction and literature review	492
10.3	Modelling seasonality in financial data	493
10.4	Estimating simple piecewise linear functions	500
10.5	Markov switching models	502
10.6	A Markov switching model for the real exchange rate	503
10.7	A Markov switching model for the gilt–equity yield ratio	506
10.8	Estimating Markov switching models in EViews	510
10.9	Threshold autoregressive models	513
10.10	Estimation of threshold autoregressive models	515
10.11	Specification tests in the context of Markov switching and threshold autoregressive models: a cautionary note	516
10.12	A SETAR model for the French franc–German mark exchange rate	517
10.13	Threshold models and the dynamics of the FTSE 100 index and index futures markets	519
10.14	A note on regime switching models and forecasting accuracy	523
11	Panel data	526
11.1	Introduction – what are panel techniques and why are they used?	526
11.2	What panel techniques are available?	528
11.3	The fixed effects model	529
11.4	Time-fixed effects models	531
11.5	Investigating banking competition using a fixed effects model	532
11.6	The random effects model	536
11.7	Panel data application to credit stability of banks in Central and Eastern Europe	537
11.8	Panel data with EViews	541
11.9	Panel unit root and cointegration tests	547
11.10	Further reading	557
12	Limited dependent variable models	559
12.1	Introduction and motivation	559
12.2	The linear probability model	560

12.3	The logit model	562
12.4	Using a logit to test the pecking order hypothesis	563
12.5	The probit model	565
12.6	Choosing between the logit and probit models	565
12.7	Estimation of limited dependent variable models	565
12.8	Goodness of fit measures for linear dependent variable models	567
12.9	Multinomial linear dependent variables	568
12.10	The pecking order hypothesis revisited – the choice between financing methods	571
12.11	Ordered response linear dependent variables models	574
12.12	Are unsolicited credit ratings biased downwards? An ordered probit analysis	574
12.13	Censored and truncated dependent variables	579
12.14	Limited dependent variable models in EViews	583
	Appendix: The maximum likelihood estimator for logit and probit models	589
13	Simulation methods	591
13.1	Motivations	591
13.2	Monte Carlo simulations	592
13.3	Variance reduction techniques	593
13.4	Bootstrapping	597
13.5	Random number generation	600
13.6	Disadvantages of the simulation approach to econometric or financial problem solving	601
13.7	An example of Monte Carlo simulation in econometrics: deriving a set of critical values for a Dickey–Fuller test	603
13.8	An example of how to simulate the price of a financial option	608
13.9	An example of bootstrapping to calculate capital risk requirements	613
14	Conducting empirical research or doing a project or dissertation in finance	626
14.1	What is an empirical research project and what is it for?	626
14.2	Selecting the topic	627
14.3	Sponsored or independent research?	629
14.4	The research proposal	631
14.5	Working papers and literature on the internet	631
14.6	Getting the data	633
14.7	Choice of computer software	634
14.8	Methodology	634
14.9	Event studies	634
14.10	Tests of the CAPM and the Fama–French Methodology	648



14.11	How might the finished project look?	662
14.12	Presentational issues	666
Appendix 1	Sources of data used in this book	667
Appendix 2	Tables of statistical distributions	668
	Glossary	680
	References	697
	Index	710

Figures

1.1	Steps involved in forming an econometric model	page 11		
2.1	A plot of hours studied (x) against grade-point average (y)	30	3.8	Effect on the standard errors of the coefficient estimates when $(x_t - \bar{x})$ are widely dispersed 96
2.2	Examples of different straight line graphs	30	3.9	Effect on the standard errors of x_t^2 large 96
2.3	Examples of quadratic functions	31	3.10	Effect on the standard errors of x_t^2 small 97
2.4	A plot of an exponential function	34	3.11	The t -distribution versus the normal 101
2.5	A plot of a logarithmic function	35	3.12	Rejection regions for a two-sided 5% hypothesis test 103
2.6	The tangent to a curve	39	3.13	Rejection region for a one-sided hypothesis test of the form $H_0 : \beta = \beta^*$, $H_1 : \beta < \beta^*$ 104
2.7	The probability distribution function for the sum of two dice	58	3.14	Rejection region for a one-sided hypothesis test of the form $H_0 : \beta = \beta^*$, $H_1 : \beta > \beta^*$ 104
2.8	The pdf for a normal distribution	59	3.15	Critical values and rejection regions for a $t_{20;5\%}$ 108
2.9	The cdf for a normal distribution	60	3.16	Frequency distribution of t -ratios of mutual fund alphas (gross of transactions costs). <i>Source</i> : Jensen (1968). Reprinted with the permission of Blackwell Publishers 114
2.10	A normal versus a skewed distribution	67	3.17	Frequency distribution of t -ratios of mutual fund alphas (net of transactions costs). <i>Source</i> : Jensen (1968). Reprinted with the permission of Blackwell Publishers 114
2.11	A normal versus a leptokurtic distribution	67	3.18	Performance of UK unit trusts, 1979–2000 116
3.1	Scatter plot of two variables, y and x	77	4.1	$R^2 = 0$ demonstrated by a flat estimated line, i.e. a zero slope coefficient 153
3.2	Scatter plot of two variables with a line of best fit chosen by eye	79	4.2	$R^2 = 1$ when all data points lie exactly on the estimated line 154
3.3	Method of OLS fitting a line to the data by minimising the sum of squared residuals	79		
3.4	Plot of a single observation, together with the line of best fit, the residual and the fitted value	80		
3.5	Scatter plot of excess returns on fund XXX versus excess returns on the market portfolio	82		
3.6	No observations close to the y -axis	84		
3.7	Effect on the standard errors of the coefficient estimates when $(x_t - \bar{x})$ are narrowly dispersed	95		

5.1	Effect of no intercept on a regression line	181	model with negative coefficient: $y_t = -0.5y_{t-1} + u_t$	272	
5.2	Graphical illustration of heteroscedasticity	182	6.7	Sample autocorrelation and partial autocorrelation functions for a non-stationary model (i.e. a unit coefficient): $y_t = y_{t-1} + u_t$	272
5.3	Plot of \hat{u}_t against \hat{u}_{t-1} , showing positive autocorrelation	191	6.8	Sample autocorrelation and partial autocorrelation functions for an ARMA(1, 1) model: $y_t = 0.5y_{t-1} + 0.5u_{t-1} + u_t$	273
5.4	Plot of \hat{u}_t over time, showing positive autocorrelation	191	6.9	Use of in-sample and out-of-sample periods for analysis	286
5.5	Plot of \hat{u}_t against \hat{u}_{t-1} , showing negative autocorrelation	192	7.1	Impulse responses and standard error bands for innovations in unexpected inflation equation errors	343
5.6	Plot of \hat{u}_t over time, showing negative autocorrelation	192	7.2	Impulse responses and standard error bands for innovations in the dividend yields	343
5.7	Plot of \hat{u}_t against \hat{u}_{t-1} , showing no autocorrelation	193	8.1	Value of R^2 for 1,000 sets of regressions of a non-stationary variable on another independent non-stationary variable	354
5.8	Plot of \hat{u}_t over time, showing no autocorrelation	193	8.2	Value of t -ratio of slope coefficient for 1,000 sets of regressions of a non-stationary variable on another independent non-stationary variable	355
5.9	Rejection and non-rejection regions for DW test	196	8.3	Example of a white noise process	358
5.10	Regression residuals from stock return data, showing large outlier for October 1987	212	8.4	Time series plot of a random walk versus a random walk with drift	359
5.11	Possible effect of an outlier on OLS estimation	213	8.5	Time series plot of a deterministic trend process	359
5.12	Plot of a variable showing suggestion for break date	231	8.6	Autoregressive processes with differing values of ϕ (0, 0.8, 1)	360
6.1	Autocorrelation function for sample MA(2) process	259	9.1	Daily S&P returns for August 2003–August 2013	423
6.2	Sample autocorrelation and partial autocorrelation functions for an MA(1) model: $y_t = -0.5u_{t-1} + u_t$	270	9.2	The problem of local optima in maximum likelihood estimation	433
6.3	Sample autocorrelation and partial autocorrelation functions for an MA(2) model: $y_t = 0.5u_{t-1} - 0.25u_{t-2} + u_t$	270	9.3	News impact curves for S&P500 returns using coefficients implied from GARCH and GJR model estimates	445
6.4	Sample autocorrelation and partial autocorrelation functions for a slowly decaying AR(1) model: $y_t = 0.9y_{t-1} + u_t$	271	9.4	Three approaches to hypothesis testing under maximum likelihood	452
6.5	Sample autocorrelation and partial autocorrelation functions for a more rapidly decaying AR(1) model: $y_t = 0.5y_{t-1} + u_t$	271			
6.6	Sample autocorrelation and partial autocorrelation functions for a more rapidly decaying AR(1)				



9.5	Time-varying hedge ratios derived from symmetric and asymmetric BEKK models for FTSE returns. <i>Source</i> : Brooks, Henry and Persaud (2002)	478		distribution with the same mean and variance. <i>Source</i> : Brooks and Persaud (2001b)	507
10.1	Sample time series plot illustrating a regime shift	491	10.6	Value of GEYR and probability that it is in the High GEYR regime for the UK. <i>Source</i> : Brooks and Persaud (2001b)	509
10.2	Use of intercept dummy variables for quarterly data	494	12.1	The fatal flaw of the linear probability model	561
10.3	Use of slope dummy variables	497	12.2	The logit model	562
10.4	Piecewise linear model with threshold x^*	501	12.3	Modelling charitable donations as a function of income	580
10.5	Unconditional distribution of US GEYR together with a normal		12.4	Fitted values from the failure probit regression	587

Tables

1.1	How to construct a series in real terms from a nominal one	page 11	5.2	Determinants and impacts of sovereign credit ratings	243
2.1	Sample data on hours of study and grades	29	5.3	Do ratings add to public information?	245
3.1	Sample data on fund XXX to motivate OLS estimation	82	5.4	What determines reactions to ratings announcements?	247
3.2	Critical values from the standard normal versus t -distribution	102	6.1	Uncovered interest parity test results	283
3.3	Classifying hypothesis testing errors and correct conclusions	110	6.2	Forecast error aggregation	292
3.4	Summary statistics for the estimated regression results for (3.34)	113	7.1	Call bid–ask spread and trading volume regression	321
3.5	Summary statistics for unit trust returns, January 1979–May 2000	115	7.2	Put bid–ask spread and trading volume regression	321
3.6	CAPM regression results for unit trust returns, January 1979–May 2000	116	7.3	Granger causality tests and implied restrictions on VAR models	335
3.7	Is there an overreaction effect in the UK stock market?	119	7.4	Marginal significance levels associated with joint F -tests	341
3.8	Part of the EViews regression output revisited	121	7.5	Variance decompositions for the property sector index residuals	342
4.1	Hedonic model of rental values in Quebec City, 1990. Dependent variable: Canadian dollars per month	157	8.1	Critical values for DF tests (Fuller, 1976, p. 373)	362
4.2	OLS and quantile regression results for the Magellan fund	165	8.2	Recursive unit root tests for interest rates allowing for structural breaks	368
4A.1	Principal component ordered eigenvalues for Dutch interest rates, 1962–70	173	8.3	DF tests on log-prices and returns for high frequency FTSE data	381
4A.2	Factor loadings of the first and second principal components for Dutch interest rates, 1962–70	174	8.4	Estimated potentially cointegrating equation and test for cointegration for high frequency FTSE data	382
5.1	Constructing a series of lagged values and first differences	190	8.5	Estimated error correction model for high frequency FTSE data	382
			8.6	Comparison of out-of-sample forecasting accuracy	383

8.7	Trading profitability of the error correction model with cost of carry	385	11.1	Tests of banking market equilibrium with fixed effects panel models	534
8.8	Cointegration tests of PPP with European data	392	11.2	Tests of competition in banking with fixed effects panel models	535
8.9	DF tests for international bond indices	393	11.3	Results of random effects panel regression for credit stability of Central and East European banks	540
8.10	Cointegration tests for pairs of international bond indices	394	11.4	Panel unit root test results for economic growth and financial development	553
8.11	Johansen tests for cointegration between international bond yields	394	11.5	Panel cointegration test results for economic growth and financial development	554
8.12	Variance decompositions for VAR of international bond yields	396	12.1	Logit estimation of the probability of external financing	564
8.13	Impulse responses for VAR of international bond yields	397	12.2	Multinomial logit estimation of the type of external financing	573
8.14	Tests of the expectations hypothesis using the US zero coupon yield curve with monthly data	400	12.3	Ordered probit model results for the determinants of credit ratings	577
9.1	GARCH versus implied volatility	457	12.4	Two-step ordered probit model allowing for selectivity bias in the determinants of credit ratings	578
9.2	EGARCH versus implied volatility	458	12.5	Marginal effects for logit and probit models for probability of MSc failure	588
9.3	Out-of-sample predictive power for weekly volatility forecasts	460	13.1	EGARCH estimates for currency futures returns	616
9.4	Comparisons of the relative information content of out-of-sample volatility forecasts	461	13.2	Autoregressive volatility estimates for currency futures returns	617
9.5	Hedging effectiveness: summary statistics for portfolio returns	477	13.3	Minimum capital risk requirements for currency futures as a percentage of the initial value of the position	620
10.1	Values and significances of days of the week coefficients	496	14.1	Journals in finance and econometrics	630
10.2	Day-of-the-week effects with the inclusion of interactive dummy variables with the risk proxy	499	14.2	Useful internet sites for financial literature	632
10.3	Estimates of the Markov switching model for real exchange rates	505	14.3	Fama and MacBeth's results on testing the CAPM	652
10.4	Estimated parameters for the Markov switching models	508	14.4	Results from Fama–MacBeth procedure using EViews	661
10.5	SETAR model for FRF–DEM	518	14.5	Suggested structure for a typical dissertation or project	662
10.6	FRF–DEM forecast accuracies	519			
10.7	Linear AR(3) model for the basis	521			
10.8	A two-threshold SETAR model for the basis	522			

Boxes

1.1	Examples of the uses of econometrics	page 2	5.2	'Solutions' for heteroscedasticity	186
1.2	Time series data	4	5.3	Conditions for <i>DW</i> to be a valid test	197
1.3	Log returns	8	5.4	Conducting a Breusch–Godfrey test	198
1.4	Points to consider when reading a published paper	13	5.5	The Cochrane–Orcutt procedure	201
1.5	Features of EViews	22	5.6	Observations for the dummy variable	212
2.1	The roots of a quadratic equation	32	5.7	Conducting a Chow test	226
2.2	Manipulating powers and their indices	33	6.1	The stationarity condition for an $AR(p)$ model	260
2.3	The laws of logs	35	6.2	The invertibility condition for an $MA(2)$ model	267
2.4	The population and the sample	62	6.3	Naive forecasting methods	288
3.1	Names for γ and x s in regression models	76	7.1	Determining whether an equation is identified	310
3.2	Reasons for the inclusion of the disturbance term	78	7.2	Conducting a Hausman test for exogeneity	312
3.3	Assumptions concerning disturbance terms and their interpretation	91	7.3	Forecasting with VARs	334
3.4	Standard error estimators	95	8.1	Stationarity tests	365
3.5	Conducting a test of significance	103	8.2	Multiple cointegrating relationships	379
3.6	Carrying out a hypothesis test using confidence intervals	106	9.1	Testing for 'ARCH effects'	426
3.7	The test of significance and confidence interval approaches compared	107	9.2	Estimating an ARCH or GARCH model	431
3.8	Type I and type II errors	111	9.3	Using maximum likelihood estimation in practice	434
3.9	Reasons for stock market overreactions	117	10.1	How do dummy variables work?	494
3.10	Ranking stocks and forming portfolios	118	11.1	Fixed or random effects?	537
3.11	Portfolio monitoring	118	12.1	Parameter interpretation for probit and logit models	566
4.1	The relationship between the regression F -statistic and R^2	158	12.2	The differences between censored and truncated dependent variables	581
4.2	Selecting between models	160	13.1	Conducting a Monte Carlo simulation	593
5.1	Conducting White's test	184	13.2	Re-sampling the data	599



13.3	Re-sampling from the residuals	600	13.6	Generating draws from a GARCH process	609
13.4	Setting up a Monte Carlo simulation	604	14.1	Possible types of research project	628
13.5	Simulating the price of an Asian option	608			

Screenshots

1.1	Creating a workfile	page 16	5.4	Chow test for parameter stability	233
1.2	Importing Excel data into the workfile	17	5.5	Plotting recursive coefficient estimates	235
1.3	The workfile containing loaded data	18	5.6	CUSUM test graph	236
1.4	Summary statistics for a series	20	6.1	Estimating the correlogram	277
1.5	A line graph	21	6.2	The options available when producing forecasts	297
2.1	Setting up a variance-covariance matrix in Excel	52	6.3	Dynamic forecasts for the percentage changes in house prices	298
2.2	The spreadsheet for constructing the efficient frontier	53	6.4	Static forecasts for the percentage changes in house prices	299
2.3	Completing the Solver window	54	6.5	Estimating exponential smoothing models	300
2.4	A plot of the completed efficient frontier	55	7.1	Estimating the inflation equation	324
2.5	The capital market line and efficient frontier	56	7.2	Estimating the rsandp equation	327
2.6	Sample summary statistics in EViews	68	7.3	VAR inputs screen	344
3.1	How to deal with dated observations in EViews	87	7.4	Constructing the VAR impulse responses	349
3.2	Summary statistics for spot and futures	88	7.5	Combined impulse response graphs	349
3.3	Equation estimation window	89	7.6	Variance decomposition graphs	350
3.4	Estimation results	90	8.1	Options menu for unit root tests	370
3.5	Plot of two series	125	8.2	Actual, fitted and residual plot to check for stationarity	401
4.1	Stepwise procedure equation estimation window	148	8.3	Johansen cointegration test	404
4.2	Stepwise procedure estimation options window	148	8.4	VAR specification for Johansen tests	409
4.3	Quantile regression estimation window	166	9.1	Estimating a GARCH-type model	436
4.4	Conducting PCA in EViews	176	9.2	GARCH model estimation options	437
5.1	Regression options window	189	9.3	Forecasting from GARCH models	450
5.2	Non-normality test results	211	9.4	Dynamic forecasts of the conditional variance	450
5.3	Regression residuals, actual values and fitted series	215	9.5	Static forecasts of the conditional variance	451
			9.6	Making a system	480



9.7	Multivariate GARCH estimation options	481	11.2	Panel workfile structure window	543
10.1	Estimating a Markov switching model	511	11.3	Panel unit root test window	556
10.2	Smoothed probabilities of being in regimes 1 and 2	513	12.1	Equation estimation window for limited dependent variables	584
11.1	Panel workfile create window	542	12.2	Equation estimation options for limited dependent variables	585
			13.1	Running an EViews program	605



Preface to the third edition

Sales of the first two editions of this book surpassed expectations (at least those of the author). Almost all of those who have contacted the author seem to like the book, and while other textbooks have been published since in the broad area of financial econometrics, none are really at the introductory level. All of the motivations for the first edition, described below, seem just as important today. Given that the book seems to have gone down well with readers, I have left the style largely unaltered but changed the structure slightly and added new material.

The main motivations for writing the first edition of the book were:

- To write a book that focused on *using and applying* the techniques rather than deriving proofs and learning formulae.
- To write an accessible textbook that required no prior knowledge of econometrics, but which also covered more recently developed approaches usually only found in more advanced texts.
- To use examples and terminology from finance rather than economics since there are many introductory texts in econometrics aimed at students of economics but none for students of finance.
- To litter the book with case studies of the use of econometrics in practice taken from the academic finance literature.
- To include sample instructions, screen dumps and computer output from a popular econometrics package. This enabled readers to see how the techniques can be implemented in practice.
- To develop a companion web site containing answers to end of chapter questions, PowerPoint slides and other supporting materials.

What is new in the third edition

The third edition includes a number of important new features:

- (1) Students of finance have enormously varying backgrounds, and in particular varying levels of training in elementary mathematics and statistics. In order to make the book more self-contained, the material that was previously buried in an appendix at the end of the book has now been considerably expanded and enhanced, and is now placed in a new chapter 2. As a result, all of the previous chapters 2 to 13 have been shunted forward by a chapter (so the

previous chapter 2 becomes chapter 3, 3 becomes 4, and so on). What was the concluding chapter in the second edition, chapter 14, has now been removed (with some of the content worked into other chapters) so that there are also fourteen chapters in the third edition.

- (2) An extensive glossary has been added at the end of the book to succinctly explain all of the technical terms used in the text.
- (3) As a result of the length of time it took to write the book, to produce the final product and the time that has elapsed since then, the data and examples used in the second edition are already several years old. The data, EViews instructions and screenshots have been fully updated. EViews version 8.0, the latest available at the time of writing, has been used throughout. The data continue to be drawn from the same freely available sources as in the previous edition.
- (4) Two of the most important uses of statistical models by students in their courses tend to be the methodology developed in a series of papers by Fama and French, and the event study approach. Both of these are now described in detail with examples in chapter 14.
- (5) New material has been added in the appropriate places in the book covering panel unit root and cointegration tests; measurement error in variables; unit root testing with structural breaks; and conditional correlation models.

Motivations for the first edition

This book had its genesis in two sets of lectures given annually by the author at the ICMA Centre (formerly ISMA Centre), Henley Business School, University of Reading and arose partly from several years of frustration at the lack of an appropriate textbook. In the past, finance was but a small sub-discipline drawn from economics and accounting, and therefore it was generally safe to assume that students of finance were well grounded in economic principles; econometrics would be taught using economic motivations and examples.

However, finance as a subject has taken on a life of its own in recent years. Drawn in by perceptions of exciting careers in the financial markets, the number of students of finance grew phenomenally all around the world. At the same time, the diversity of educational backgrounds of students taking finance courses has also expanded. It is not uncommon to find undergraduate students of finance even without advanced high-school qualifications in mathematics or economics. Conversely, many with PhDs in physics or engineering are also attracted to study finance at the Masters level. Unfortunately, authors of textbooks failed to keep pace with the change in the nature of students. In my opinion, the currently available textbooks fall short of the requirements of this market in three main regards, which this book seeks to address:

- (1) Books fall into two distinct and non-overlapping categories: the introductory and the advanced. Introductory textbooks are at the appropriate level for students with limited backgrounds in mathematics or statistics, but their focus is too narrow. They often spend too long deriving the most basic results, and

treatment of important, interesting and relevant topics (such as simulations methods, VAR modelling, etc.) is covered in only the last few pages, if at all. The more advanced textbooks, meanwhile, usually require a quantum leap in the level of mathematical ability assumed of readers, so that such books cannot be used on courses lasting only one or two semesters, or where students have differing backgrounds. In this book, I have tried to sweep a broad brush over a large number of different econometric techniques that are relevant to the analysis of financial and other data.

- (2) Many of the currently available textbooks with broad coverage are too theoretical in nature and students can often, after reading such a book, still have no idea of how to tackle real-world problems themselves, even if they have mastered the techniques in theory. To this end, in this book, I have tried to present examples of the use of the techniques in finance, together with annotated computer instructions and sample outputs for an econometrics package (EViews). This should assist students who wish to learn how to estimate models for themselves – for example, if they are required to complete a project or dissertation. Some examples have been developed especially for this book, while many others are drawn from the academic finance literature. In my opinion, this is an essential but rare feature of a textbook that should help to show students how econometrics is really applied. It is also hoped that this approach will encourage some students to delve deeper into the literature, and will give useful pointers and stimulate ideas for research projects. It should, however, be stated at the outset that the purpose of including examples from the academic finance print is not to provide a comprehensive overview of the literature or to discuss all of the relevant work in those areas, but rather to illustrate the techniques. Therefore, the literature reviews may be considered deliberately deficient, with interested readers directed to the suggested readings and the references therein.
- (3) With few exceptions, almost all textbooks that are aimed at the introductory level draw their motivations and examples from economics, which may be of limited interest to students of finance or business. To see this, try motivating regression relationships using an example such as the effect of changes in income on consumption and watch your audience, who are primarily interested in business and finance applications, slip away and lose interest in the first ten minutes of your course.

Who should read this book?

The intended audience is undergraduates or Masters/MBA students who require a broad knowledge of modern econometric techniques commonly employed in the finance literature. It is hoped that the book will also be useful for researchers (both academics and practitioners), who require an introduction to the statistical tools commonly employed in the area of finance. The book can be used for courses covering financial time-series analysis or financial econometrics in undergraduate or postgraduate programmes in finance, financial economics, securities and investments.



Although the applications and motivations for model-building given in the book are drawn from finance, the empirical testing of theories in many other disciplines, such as management studies, business studies, real estate, economics and so on, may usefully employ econometric analysis. For this group, the book may also prove useful.

Finally, while the present text is designed mainly for students at the undergraduate or Masters level, it could also provide introductory reading in financial time series modelling for finance doctoral programmes where students have backgrounds which do not include courses in modern econometric techniques.

Pre-requisites for good understanding of this material

In order to make the book as accessible as possible, no prior knowledge of statistics, econometrics or algebra is required, although those with a prior exposure to calculus, algebra (including matrices) and basic statistics will be able to progress more quickly. The emphasis throughout the book is on a valid application of the techniques to real data and problems in finance.

In the finance and investment area, it is assumed that the reader has knowledge of the fundamentals of corporate finance, financial markets and investment. Therefore, subjects such as portfolio theory, the capital asset pricing model (CAPM) and Arbitrage Pricing Theory (APT), the efficient markets hypothesis, the pricing of derivative securities and the term structure of interest rates, which are frequently referred to throughout the book, are not explained from first principles in this text. There are very many good books available in corporate finance, in investments and in futures and options, including those by Brealey and Myers (2013), Bodie, Kane and Marcus (2011) and Hull (2011) respectively.



Acknowledgements

I am grateful to Gita Persand, Olan Henry, James Chong and Apostolos Katsaris, who assisted with various parts of the software applications for the first edition. I am also grateful to Hilary Feltham for assistance with chapter 2 and to Simone Varotto for useful discussions and advice concerning the EViews example used in chapter 11.

I would also like to thank Simon Burke, James Chong and Con Keating for detailed and constructive comments on various drafts of the first edition, Simon Burke for suggestions on parts of the second edition and Jo Cox for comments on part of the third edition. The first and second editions additionally benefited from the comments, suggestions and questions of Peter Burr ridge, Kyong-wook Choi, Rishi Chopra, Araceli Ortega Diaz, Xiaoming Ding, Thomas Eilertsen, Waleid Eldien, Andrea Gheno, Christopher Gilbert, Kimon Gomozi as, Cherif Guer mat, Abid Hameed, Arty Khemlani, Margaret Lynch, David McCaffrey, Tehri Jokipii, Emese Lazar, Zhao Liuyan, Dimitri Lvov, Bill McCabe, Junshi Ma, David Merchan, Victor Murinde, Mikael Petitjean, Marcelo Perlin, Thai Pham, Jean-Sebastien Pourchet, Marcel Prokopczuk, Guilherme Silva, Jerry Sin, Silvia Stanescu, Yiguo Sun, Li Qui, Panagiotis Varlagas, Jakub Vojtek, Jue Wang and Meng-Feng Yen.

A number of people sent useful e-mails pointing out typos or inaccuracies in the first edition. To this end, I am grateful to Merlyn Foo, Jan de Gooijer and his colleagues, Mikael Petitjean, Fred Sterbenz and Birgit Strikholm.

Useful comments and software support from Quantitative Micro Software (QMS) (now IHS Global) are gratefully acknowledged. Any remaining errors are mine alone.

The publisher and author have used their best endeavours to ensure that the URLs for external web sites referred to in this book are correct and active at the time of going to press. However, the publisher and author have no responsibility for the web sites and can make no guarantee that a site will remain live or that the content is or will remain appropriate.

1

Introduction

Learning econometrics is in many ways like learning a new language. To begin with, nothing makes sense and it is as if it is impossible to see through the fog created by all the unfamiliar terminology. While the way of writing the models – the *notation* – may make the situation appear more complex, in fact it is supposed to achieve the exact opposite. The ideas themselves are mostly not so complicated, it is just a matter of learning enough of the language that everything fits into place. So if you have never studied the subject before, then persevere through this preliminary chapter and you will hopefully be on your way to being fully fluent in econometrics!

Learning outcomes

In this chapter, you will learn how to

- Compare nominal and real series and convert one to the other
- Distinguish between different types of data
- Describe the key steps involved in building an econometric model
- Calculate asset price returns
- Deflate series to allow for inflation
- Construct a workfile, import data and accomplish simple tasks in EViews

The chapter sets the scene for the book by discussing in broad terms the questions of what econometrics is, and what are the ‘stylised facts’ describing financial data that researchers in this area typically try to capture in their models? Some discussion is presented on the kinds of data we encounter in finance and how to work with them. Finally, the chapter collects together a number of preliminary issues relating to the construction of econometric models in finance and introduces the software that will be used in the remainder of the book for estimating the models.

Box 1.1 Examples of the uses of econometrics

- (1) Testing whether financial markets are weak-form informationally efficient
- (2) Testing whether the capital asset pricing model (CAPM) or arbitrage pricing theory (APT) represent superior models for the determination of returns on risky assets
- (3) Measuring and forecasting the volatility of bond returns
- (4) Explaining the determinants of bond credit ratings used by the ratings agencies
- (5) Modelling long-term relationships between prices and exchange rates
- (6) Determining the optimal hedge ratio for a spot position in oil
- (7) Testing technical trading rules to determine which makes the most money
- (8) Testing the hypothesis that earnings or dividend announcements have no effect on stock prices
- (9) Testing whether spot or futures markets react more rapidly to news
- (10) Forecasting the correlation between the stock indices of two countries.

1.1

What is econometrics?

The literal meaning of the word econometrics is ‘measurement in economics’. The first four letters of the word suggest correctly that the origins of econometrics are rooted in economics. However, the main techniques employed for studying economic problems are of equal importance in financial applications. As the term is used in this book, financial econometrics will be defined as the *application of statistical techniques to problems in finance*. Financial econometrics can be useful for testing theories in finance, determining asset prices or returns, testing hypotheses concerning the relationships between variables, examining the effect on financial markets of changes in economic conditions, forecasting future values of financial variables and for financial decision-making. A list of possible examples of where econometrics may be useful is given in box 1.1.

The list in box 1.1 is of course by no means exhaustive, but it hopefully gives some flavour of the usefulness of econometric tools in terms of their financial applicability.

1.2

Is financial econometrics different from ‘economic econometrics’?

As previously stated, the tools commonly used in financial applications are fundamentally the same as those used in economic applications, although the emphasis and the sets of problems that are likely to be encountered when analysing the two

sets of data are somewhat different. Financial data often differ from macroeconomic data in terms of their frequency, accuracy, seasonality and other properties.

In economics, a serious problem is often a *lack of data at hand* for testing the theory or hypothesis of interest – this is often called a ‘small samples problem’. It might be, for example, that data are required on government budget deficits, or population figures, which are measured only on an annual basis. If the methods used to measure these quantities changed a quarter of a century ago, then only at most twenty-five of these annual observations are usefully available.

Two other problems that are often encountered in conducting applied econometric work in the arena of economics are those of *measurement error* and *data revisions*. These difficulties are simply that the data may be estimated, or measured with error, and will often be subject to several vintages of subsequent revisions. For example, a researcher may estimate an economic model of the effect on national output of investment in computer technology using a set of published data, only to find that the data for the last two years have been revised substantially in the next, updated publication.

These issues are usually of less concern in finance. Financial data come in many shapes and forms, but in general the prices and other entities that are recorded are those at which trades *actually took place*, or which were *quoted* on the screens of information providers. There exists, of course, the possibility for typos or for the data measurement method to change (for example, owing to stock index re-balancing or re-basing). But in general the measurement error and revisions problems are far less serious in the financial context.

Similarly, some sets of financial data are observed at much *higher frequencies* than macroeconomic data. Asset prices or yields are often available at daily, hourly or minute-by-minute frequencies. Thus the number of observations available for analysis can potentially be very large – perhaps thousands or even millions, making financial data the envy of macro-econometricians! The implication is that more powerful techniques can often be applied to financial than economic data, and that researchers may also have more confidence in the results.

Furthermore, the analysis of financial data also brings with it a number of new problems. While the difficulties associated with handling and processing such a large amount of data are not usually an issue given recent and continuing advances in computer power, financial data often have a number of additional characteristics. For example, financial data are often considered very ‘noisy’, which means that it is more difficult to separate *underlying trends or patterns* from random and uninteresting features. Financial data are also almost always not normally distributed in spite of the fact that most techniques in econometrics assume that they are. High frequency data often contain additional ‘patterns’ which are the result of the way that the market works, or the way that prices are recorded. These features need to be considered in the model-building process, even if they are not directly of interest to the researcher.

One of the most rapidly evolving areas of financial application of statistical tools is in the modelling of market microstructure problems. ‘Market microstructure’ may broadly be defined as the process whereby *investors’ preferences and desires are translated into financial market transactions*. It is evident that microstructure effects

Box 1.2 Time series data

<i>Series</i>	<i>Frequency</i>
Industrial production	Monthly or quarterly
Government budget deficit	Annually
Money supply	Weekly
The value of a stock	As transactions occur

are important and represent a key difference between financial and other types of data. These effects can potentially impact on many other areas of finance. For example, market rigidities or frictions can imply that current asset prices do not fully reflect future expected cashflows (see the discussion in chapter 10 of this book). Also, investors are likely to require compensation for holding securities that are illiquid, and therefore embody a risk that they will be difficult to sell owing to the relatively high probability of a lack of willing purchasers at the time of desired sale. Measures such as volume or the time between trades are sometimes used as proxies for market liquidity.

A comprehensive survey of the literature on market microstructure is given by Madhavan (2000). He identifies several aspects of the market microstructure literature, including price formation and price discovery, issues relating to market structure and design, information and disclosure. There are also relevant books by O'Hara (1995), Harris (2002) and Hasbrouck (2007). At the same time, there has been considerable advancement in the sophistication of econometric models applied to microstructure problems. For example, an important innovation was the autoregressive conditional duration (ACD) model attributed to Engle and Russell (1998). An interesting application can be found in Dufour and Engle (2000), who examine the effect of the time between trades on the price-impact of the trade and the speed of price adjustment.

1.3 Types of data

There are broadly three types of data that can be employed in quantitative analysis of financial problems: time series data, cross-sectional data and panel data.

1.3.1 Time series data

Time series data, as the name suggests, are data that have been collected over a period of time on one or more variables. Time series data have associated with them a particular frequency of observation or frequency of collection of data points. The frequency is simply a measure of the *interval over*, or the *regularity with which*, the data are collected or recorded. Box 1.2 shows some examples of time series data.

A word on ‘As transactions occur’ is necessary. Much financial data does not start its life as being *regularly spaced*. For example, the price of common stock for a given company might be recorded to have changed whenever there is a new trade or quotation placed by the financial information recorder. Such recordings are very unlikely to be evenly distributed over time – for example, there may be no activity between, say, 5 p.m. when the market closes and 8.30 a.m. the next day when it reopens; there is also typically less activity around the opening and closing of the market, and around lunch time. Although there are a number of ways to deal with this issue, a common and simple approach is simply to select an appropriate frequency, and use as the observation for that time period the last prevailing price during the interval.

It is also generally a requirement that all data used in a model be of the *same frequency of observation*. So, for example, regressions that seek to estimate an arbitrage pricing model using monthly observations on macroeconomic factors must also use monthly observations on stock returns, even if daily or weekly observations on the latter are available.

The data may be *quantitative* (e.g. exchange rates, prices, number of shares outstanding), or *qualitative* (e.g. the day of the week, a survey of the financial products purchased by private individuals over a period of time, a credit rating, etc.).

Problems that could be tackled using time series data:

- How the value of a country’s stock index has varied with that country’s macroeconomic fundamentals
- How the value of a company’s stock price has varied when it announced the value of its dividend payment
- The effect on a country’s exchange rate of an increase in its trade deficit.

In all of the above cases, it is clearly the time dimension which is the most important, and the analysis will be conducted using the values of the variables over time.

1.3.2 Cross-sectional data

Cross-sectional data are data on one or more variables collected at a single point in time. For example, the data might be on:

- A poll of usage of internet stockbroking services
- A cross-section of stock returns on the New York Stock Exchange (NYSE)
- A sample of bond credit ratings for UK banks.

Problems that could be tackled using cross-sectional data:

- The relationship between company size and the return to investing in its shares
- The relationship between a country’s GDP level and the probability that the government will default on its sovereign debt.



1.3.3 Panel data

Panel data have the dimensions of both time series and cross-sections, e.g. the daily prices of a number of blue chip stocks over two years. The estimation of panel regressions is an interesting and developing area, and will be examined in detail in chapter 11.

Fortunately, virtually all of the standard techniques and analysis in econometrics are equally valid for time series and cross-sectional data. For time series data, it is usual to denote the individual observation numbers using the index t , and the total number of observations available for analysis by T . For cross-sectional data, the individual observation numbers are indicated using the index i , and the total number of observations available for analysis by N . Note that there is, in contrast to the time series case, no natural ordering of the observations in a cross-sectional sample. For example, the observations i might be on the price of bonds of different firms at a particular point in time, ordered alphabetically by company name. So, in the case of cross-sectional data, there is unlikely to be any useful information contained in the fact that Barclays follows Banco Santander in a sample of bank credit ratings, since it is purely by chance that their names both begin with the letter 'B'. On the other hand, in a time series context, the ordering of the data is relevant since the data are usually ordered chronologically.

In this book, the total number of observations in the sample will be given by T even in the context of regression equations that could apply either to cross-sectional or to time series data.

1.3.4 Continuous and discrete data

As well as classifying data as being of the time series or cross-sectional type, we could also distinguish them as being either continuous or discrete, exactly as their labels would suggest. *Continuous* data can take on any value and are not confined to take specific numbers; their values are limited only by precision. For example, the rental yield on a property could be 6.2%, 6.24% or 6.238%, and so on. On the other hand, *discrete* data can only take on certain values, which are usually integers (whole numbers), and are often defined to be count numbers.¹ For instance, the number of people in a particular underground carriage or the number of shares traded during a day. In these cases, having 86.3 passengers in the carriage or 5857 $\frac{1}{2}$ shares traded would not make sense. The simplest example of a discrete variable is a *Bernoulli* or binary random variable, which can only take the values 0 or 1 – for example, if we repeatedly tossed a coin, we could denote a head by 0 and a tail by 1.

¹ Discretely measured data do not necessarily have to be integers. For example, until they became 'decimalised', many financial asset prices were quoted to the nearest 1/16 or 1/32 of a dollar.

1.3.5 Cardinal, ordinal and nominal numbers

Another way in which we could classify numbers is according to whether they are cardinal, ordinal or nominal. *Cardinal* numbers are those where the actual numerical values that a particular variable takes have meaning, and where there is an equal distance between the numerical values. On the other hand, *ordinal* numbers can only be interpreted as providing a position or an ordering. Thus, for cardinal numbers, a figure of 12 implies a measure that is ‘twice as good’ as a figure of 6. Examples of cardinal numbers would be the price of a share or of a building, and the number of houses in a street. On the other hand, for an ordinal scale, a figure of 12 may be viewed as ‘better’ than a figure of 6, but could not be considered twice as good. Examples of ordinal numbers would be the position of a runner in a race (e.g. second place is better than fourth place, but it would make little sense to say it is ‘twice as good’) or the level reached in a computer game.

The final type of data that could be encountered would be where there is no natural ordering of the values at all, so a figure of 12 is simply different to that of a figure of 6, but could not be considered to be better or worse in any sense. Such data often arise when numerical values are arbitrarily assigned, such as telephone numbers or when codings are assigned to qualitative data (e.g. when describing the exchange that a US stock is traded on, ‘1’ might be used to denote the NYSE, ‘2’ to denote the NASDAQ and ‘3’ to denote the AMEX). Sometimes, such variables are called *nominal* variables. Cardinal, ordinal and nominal variables may require different modelling approaches or at least different treatments, as should become evident in the subsequent chapters.

1.4 Returns in financial modelling

In many of the problems of interest in finance, the starting point is a time series of prices – for example, the prices of shares in Ford, taken at 4 p.m. each day for 200 days. For a number of statistical reasons, it is preferable not to work directly with the price series, so that raw price series are usually converted into series of returns. Additionally, returns have the added benefit that they are unit-free. So, for example, if an annualised return were 10%, then investors know that they would have got back £110 for a £100 investment, or £1,100 for a £1,000 investment, and so on.

There are two methods used to calculate returns from a series of prices, and these involve the formation of simple returns, and continuously compounded returns, which are achieved as follows:

Simple returns

$$R_t = \frac{p_t - p_{t-1}}{p_{t-1}} \times 100\% \quad (1.1)$$

Continuously compounded returns

$$r_t = 100\% \times \ln \left(\frac{p_t}{p_{t-1}} \right) \quad (1.2)$$

Box 1.3 Log returns

- (1) Log-returns have the nice property that they can be interpreted as *continuously compounded returns* – so that the frequency of compounding of the return does not matter and thus returns across assets can more easily be compared.
- (2) Continuously compounded returns are *time-additive*. For example, suppose that a weekly returns series is required and daily log returns have been calculated for five days, numbered 1 to 5, representing the returns on Monday through Friday. It is valid to simply add up the five daily returns to obtain the return for the whole week:

<i>Monday return</i>	$r_1 = \ln(p_1/p_0) = \ln p_1 - \ln p_0$
<i>Tuesday return</i>	$r_2 = \ln(p_2/p_1) = \ln p_2 - \ln p_1$
<i>Wednesday return</i>	$r_3 = \ln(p_3/p_2) = \ln p_3 - \ln p_2$
<i>Thursday return</i>	$r_4 = \ln(p_4/p_3) = \ln p_4 - \ln p_3$
<i>Friday return</i>	$r_5 = \ln(p_5/p_4) = \ln p_5 - \ln p_4$
 <i>Return over the week</i>	<hr style="width: 50%; margin: 0 auto;"/> $\ln p_5 - \ln p_0 = \ln(p_5/p_0)$

where: R_t denotes the simple return at time t , r_t denotes the continuously compounded return at time t , p_t denotes the asset price at time t and \ln denotes the natural logarithm.

If the asset under consideration is a stock or portfolio of stocks, the total return to holding it is the sum of the capital gain and any dividends paid during the holding period. However, researchers often ignore any dividend payments. This is unfortunate, and will lead to an underestimation of the total returns that accrue to investors. This is likely to be negligible for very short holding periods, but will have a severe impact on cumulative returns over investment horizons of several years. Ignoring dividends will also have a distortionary effect on the cross-section of stock returns. For example, ignoring dividends will imply that ‘growth’ stocks with large capital gains will be inappropriately favoured over income stocks (e.g. utilities and mature industries) that pay high dividends.

Alternatively, it is possible to adjust a stock price time series so that the dividends are added back to generate a *total return index*. If p_t were a total return index, returns generated using either of the two formulae presented above thus provide a measure of the total return that would accrue to a holder of the asset during time t .

The academic finance literature generally employs the log-return formulation (also known as log-price relatives since they are the log of the ratio of this period’s price to the previous period’s price). Box 1.3 shows two key reasons for this.

There is, however, also a disadvantage of using the log-returns. The simple return on a portfolio of assets is a weighted average of the simple returns on the

individual assets:

$$R_{pt} = \sum_{i=1}^N w_i R_{it} \quad (1.3)$$

But this does not work for the continuously compounded returns, so that they are not additive across a portfolio. The fundamental reason why this is the case is that the log of a sum is not the same as the sum of a log, since the operation of taking a log constitutes a *non-linear transformation*. Calculating portfolio returns in this context must be conducted by first estimating the value of the portfolio at each time period and then determining the returns from the aggregate portfolio values. Or alternatively, if we assume that the asset is purchased at time $t - K$ for price p_{t-K} and then sold K periods later at price p_t , then if we calculate simple returns for each period, R_t, R_{t+1}, \dots, R_K , the aggregate return over all K periods is

$$\begin{aligned} R_{Kt} &= \frac{p_t - p_{t-K}}{p_{t-K}} = \frac{p_t}{p_{t-K}} - 1 = \left[\frac{p_t}{p_{t-1}} \times \frac{p_{t-1}}{p_{t-2}} \times \dots \times \frac{p_{t-K+1}}{p_{t-K}} \right] - 1 \\ &= [(1 + R_t)(1 + R_{t-1}) \dots (1 + R_{t-K+1})] - 1 \end{aligned} \quad (1.4)$$

In the limit, as the frequency of the sampling of the data is increased so that they are measured over a smaller and smaller time interval, the simple and continuously compounded returns will be identical.

1.4.1 Real versus nominal series and deflating nominal series

If a newspaper headline suggests that ‘house prices are growing at their fastest rate for more than a decade. A typical 3-bedroom house is now selling for £180,000, whereas in 1990 the figure was £120,000’, it is important to appreciate that this figure is almost certainly in *nominal* terms. That is, the article is referring to the actual prices of houses that existed at those points in time. The general level of prices in most economies around the world has a general tendency to rise almost all of the time, so we need to ensure that we compare prices on a like-for-like basis. We could think of part of the rise in house prices being attributable to an increase in demand for housing, and part simply arising because the prices of all goods and services are rising together. It would be useful to be able to separate the two effects, and to be able to answer the question, ‘how much have house prices risen when we remove the effects of general inflation?’ or equivalently, ‘how much are houses worth now if we measure their values in 1990-terms?’ We can do this by *deflating* the nominal house price series to create a series of *real* house prices, which is then said to be in *inflation-adjusted terms* or *at constant prices*.

Deflating a series is very easy indeed to achieve: all that is required (apart from the series to deflate) is a *price deflator series*, which is a series measuring general price levels in the economy. Series like the consumer price index (CPI), producer price index (PPI) or the GDP Implicit Price Deflator, are often used. A more detailed discussion of which is the most relevant general price index to use is beyond the