

**AN INTRODUCTION
TO
HIGH-FREQUENCY
FINANCE**

AN INTRODUCTION TO HIGH-FREQUENCY FINANCE

Michel M. Dacorogna

Zurich Re, Switzerland

Ramazan Gençay

University of Windsor, Canada

Olsen & Associates, Switzerland

Ulrich A. Müller

Olsen & Associates, Switzerland

Richard B. Olsen

Olsen & Associates, Switzerland

Olivier V. Pictet

Dynamic Asset Management, Switzerland



ACADEMIC PRESS

A Harcourt Science and Technology Company

San Diego San Francisco New York Boston London Sydney Tokyo

Cover images © 2001 PhotoDisc, Inc.

This book is printed on acid-free paper.

Copyright © 2001 by ACADEMIC PRESS

All Rights Reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the publisher.

Requests for permission to make copies of any part of the work should be mailed to:
Permissions Department, Harcourt Inc., 6277 Sea Harbor Drive,
Orlando, Florida 32887-6777

Academic Press
A Harcourt Science and Technology Company
525 B Street, Suite 1900, San Diego, California 92101-4495, USA
<http://www.academicpress.com>

Academic Press
Harcourt Place, 32 Jamestown Road, London NW1 7BY, UK
<http://www.academicpress.com>

Library of Congress Catalog Card Number: 2001088178

International Standard Book Number: 0-12-279671-3

PRINTED IN THE UNITED STATES OF AMERICA

01 02 03 04 05 06 EB 9 8 7 6 5 4 3 2 1

To our parents and families

CONTENTS

LIST OF FIGURES xv

LIST OF TABLES xix

PREFACE xxi

ACKNOWLEDGMENTS xxiii

I

INTRODUCTION

| | | |
|------------|--------------------------------------------|---|
| 1.1 | Markets: The Source of High-Frequency Data | 1 |
| 1.2 | Methodology of High-Frequency Research | 2 |
| 1.3 | Data Frequency and Market Information | 3 |
| 1.4 | New Levels of Significance | 6 |
| 1.5 | Interrelating Different Time Scales | 8 |

2

MARKETS AND DATA

| | | |
|--------------|-------------------------------------------------|----|
| 2.1 | General Remarks on Markets and Data Types | 10 |
| 2.1.1 | Spot Markets | 11 |
| 2.1.2 | Futures Markets | 12 |
| 2.1.3 | Option Markets | 13 |
| 2.2 | Foreign Exchange Markets | 13 |
| 2.2.1 | Structure of the Foreign Exchange Spot Market | 15 |
| 2.2.2 | Synthetic Cross Rates | 19 |
| 2.2.3 | Multiple Contributor Effects | 19 |
| 2.3 | Over-the-Counter Interest Rate Markets | 20 |
| 2.3.1 | Spot Interest Rates | 21 |
| 2.3.2 | Foreign Exchange Forward Rates | 22 |
| 2.4 | Interest Rate Futures | 23 |
| 2.4.1 | General Description of Interest Rate Futures | 23 |
| 2.4.2 | Implied Forward Interest Rates and Yield Curves | 25 |
| 2.5 | Bond Futures Markets | 28 |
| 2.5.1 | Bonds and Bond Futures | 28 |
| 2.5.2 | Rollover Schemes | 29 |
| 2.6 | Commodity Futures | 31 |
| 2.7 | Equity Markets | 32 |

3

TIME SERIES OF INTEREST

| | | |
|--------------|-----------------------------------------|----|
| 3.1 | Time Series and Operators | 34 |
| 3.2 | Variables in Homogeneous Time Series | 37 |
| 3.2.1 | Interpolation | 37 |
| 3.2.2 | Price | 38 |
| 3.2.3 | Return | 40 |
| 3.2.4 | Realized Volatility | 41 |
| 3.2.5 | Bid-Ask Spread | 45 |
| 3.2.6 | Tick Frequency | 46 |
| 3.2.7 | Other Variables | 46 |
| 3.2.8 | Overlapping Returns | 47 |
| 3.3 | Convolution Operators | 51 |
| 3.3.1 | Notation Used for Time Series Operators | 53 |
| 3.3.2 | Linear Operator and Kernels | 54 |
| 3.3.3 | Build-Up Time Interval | 56 |
| 3.3.4 | Homogeneous Operators and Robustness | 58 |
| 3.3.5 | Exponential Moving Average (EMA) | 59 |
| 3.3.6 | The Iterated EMA Operator | 59 |
| 3.3.7 | Moving Average (MA) | 61 |

| | | |
|--------|---------------------------------------------------------|----|
| 3.3.8 | Moving Norm, Variance, and Standard Deviation | 63 |
| 3.3.9 | Differential | 64 |
| 3.3.10 | Derivative and γ -Derivative | 66 |
| 3.3.11 | Volatility | 68 |
| 3.3.12 | Standardized Time Series, Moving Skewness, and Kurtosis | 71 |
| 3.3.13 | Moving Correlation | 71 |
| 3.3.14 | Windowed Fourier Transform | 74 |
| 3.4 | Microscopic Operators | 76 |
| 3.4.1 | Backward Shift and Time Translation Operators | 77 |
| 3.4.2 | Regular Time Series Operator | 77 |
| 3.4.3 | Microscopic Return, Difference, and Derivative | 78 |
| 3.4.4 | Microscopic Volatility | 79 |
| 3.4.5 | Tick Frequency and Activity | 79 |

4

ADAPTIVE DATA CLEANING

| | | |
|-------|-------------------------------------------------------|-----|
| 4.1 | Introduction: Using a Filter to Clean the Data | 82 |
| 4.2 | Data and Data Errors | 84 |
| 4.2.1 | Time Series of Ticks | 84 |
| 4.2.2 | Data Error Types | 85 |
| 4.3 | General Overview of the Filter | 86 |
| 4.3.1 | The Functionality of the Filter | 86 |
| 4.3.2 | Overview of the Filtering Algorithm and Its Structure | 88 |
| 4.4 | Basic Filtering Elements and Operations | 88 |
| 4.4.1 | Credibility and Trust Capital | 89 |
| 4.4.2 | Filtering of Single Scalar Quotes: The Level Filter | 91 |
| 4.4.3 | Pair Filtering: The Credibility of Returns | 93 |
| 4.4.4 | Computing the Expected Volatility | 96 |
| 4.4.5 | Pair Filtering: Comparing Quote Origins | 98 |
| 4.4.6 | A Time Scale for Filtering | 100 |
| 4.5 | The Scalar Filtering Window | 103 |
| 4.5.1 | Entering a New Quote in the Scalar Filtering Window | 104 |
| 4.5.2 | The Trust Capital of a New Scalar Quote | 104 |
| 4.5.3 | Updating the Scalar Window | 106 |
| 4.5.4 | Dismissing Quotes from the Scalar Window | 107 |
| 4.5.5 | Updating the Statistics with Credible Scalar Quotes | 108 |
| 4.5.6 | A Second Scalar Window for Old Valid Quotes | 108 |
| 4.6 | The Full-Quote Filtering Window | 109 |
| 4.6.1 | Quote Splitting Depending on the Instrument Type | 110 |
| 4.6.2 | The Basic Validity Test | 110 |
| 4.6.3 | Transforming the Filtered Variable | 112 |
| 4.7 | Univariate Filtering | 113 |

| | | |
|-------|-----------------------------------------------|-----|
| 4.7.1 | The Results of Univariate Filtering | 114 |
| 4.7.2 | Filtering in Historical and Real-Time Modes | 115 |
| 4.7.3 | Choosing the Filter Parameters | 116 |
| 4.8 | Special Filter Elements | 116 |
| 4.8.1 | Multivariate Filtering: Filtering Sparse Data | 116 |
| 4.9 | Behavior and Effects of the Data Filter | 118 |

5

BASIC STYLIZED FACTS

| | | |
|-------|------------------------------------------------------------------|-----|
| 5.1 | Introduction | 121 |
| 5.2 | Price Formation Process | 123 |
| 5.2.1 | Negative First-Order Autocorrelation of Returns | 123 |
| 5.2.2 | Discreteness of Quoted Spreads | 125 |
| 5.2.3 | Short-Term Triangular Arbitrage | 127 |
| 5.3 | Institutional Structure and Exogenous Impacts | 127 |
| 5.3.1 | Institutional Framework | 127 |
| 5.3.2 | Positive Impact of Official Interventions | 129 |
| 5.3.3 | Mixed Effect of News | 129 |
| 5.4 | Distributional Properties of Returns | 132 |
| 5.4.1 | Finite Variance, Symmetry and Decreasing Fat-Tailedness | 132 |
| 5.4.2 | The Tail Index of Return Distributions | 135 |
| 5.4.3 | Extreme Risks in Financial Markets | 144 |
| 5.5 | Scaling Laws | 147 |
| 5.5.1 | Empirical Evidence | 147 |
| 5.5.2 | Distributions and Scaling Laws | 151 |
| 5.5.3 | A Simple Model of the Market Maker Bias | 154 |
| 5.5.4 | Limitations of the Scaling Laws | 158 |
| 5.6 | Autocorrelation and Seasonality | 160 |
| 5.6.1 | Autocorrelations of Returns and Volatility | 161 |
| 5.6.2 | Seasonal Volatility: Across Markets for OTC Instruments | 163 |
| 5.6.3 | Seasonal Volatility: U-Shaped for Exchange Traded Instruments | 167 |
| 5.6.4 | Deterministic Volatility in Eurofutures Contracts | 169 |
| 5.6.5 | Bid-Ask Spreads | 170 |

6

MODELING SEASONAL VOLATILITY

| | | |
|-------|-------------------------------------------------------------|-----|
| 6.1 | Introduction | 174 |
| 6.2 | A Model of Market Activity | 175 |
| 6.2.1 | Seasonal Patterns of the Volatility and Presence of Markets | 175 |

| | | |
|-------|------------------------------------------------------------------------------------------|-----|
| 6.2.2 | Modeling the Volatility Patterns with an Alternative Time Scale and an Activity Variable | 176 |
| 6.2.3 | Market Activity and Scaling Law | 177 |
| 6.2.4 | Geographical Components of Market Activity | 178 |
| 6.2.5 | A Model of Intraday Market Activity | 179 |
| 6.2.6 | Interpretation of the Activity Modeling Results | 183 |
| 6.3 | A New Business Time Scale (ϑ -Scale) | 188 |
| 6.3.1 | Definition of the ϑ -Scale | 188 |
| 6.3.2 | Adjustments of the ϑ -Scale Definition | 189 |
| 6.3.3 | A Ratio Test for the ϑ -Scale Quality | 192 |
| 6.4 | Filtering Intraday Seasonalities with Wavelets | 193 |

7

REALIZED VOLATILITY DYNAMICS

| | | |
|-------|----------------------------------------------------|-----|
| 7.1 | Introduction | 197 |
| 7.2 | The Bias of Realized Volatility and Its Correction | 198 |
| 7.3 | Conditional Heteroskedasticity | 204 |
| 7.3.1 | Autocorrelation of Volatility in ϑ -Time | 204 |
| 7.3.2 | Short and Long Memory | 207 |
| 7.4 | The Heterogeneous Market Hypothesis | 209 |
| 7.4.1 | Volatilities of Different Time Resolutions | 210 |
| 7.4.2 | Asymmetric Lead-Lag Correlation of Volatilities | 211 |
| 7.4.3 | Conditional Predictability | 215 |

8

VOLATILITY PROCESSES

| | | |
|-------|----------------------------------------------------|-----|
| 8.1 | Introduction | 219 |
| 8.2 | Intraday Volatility and GARCH Models | 221 |
| 8.2.1 | Parameter Estimation of GARCH Models | 222 |
| 8.2.2 | Temporal Aggregation of GARCH Models | 224 |
| 8.2.3 | Estimates of GARCH(1,1) for Various Frequencies | 226 |
| 8.3 | Modeling Heterogeneous Volatilities | 231 |
| 8.3.1 | The HARCH Model | 231 |
| 8.3.2 | HARCH and Market Components | 234 |
| 8.3.3 | Generalization of the Process Equation | 237 |
| 8.3.4 | EMA-HARCH Model | 237 |
| 8.3.5 | Estimating HARCH and EMA-HARCH Models | 239 |
| 8.3.6 | HARCH in Interest Rate Modeling | 242 |
| 8.4 | Forecasting Short-Term Volatility | 243 |
| 8.4.1 | A Framework to Measure the Forecasting Performance | 243 |
| 8.4.2 | Performance of ARCH-Type Models | 246 |

9

FORECASTING RISK AND RETURN

| | | |
|--------------|-------------------------------------------------|-----|
| 9.1 | Introduction to Forecasting | 248 |
| 9.2 | Forecasting Volatility for Value-at-Risk | 250 |
| 9.2.1 | Three Simple Volatility Forecasting Models | 250 |
| 9.2.2 | Choosing the Best Volatility Forecasting Model | 254 |
| 9.3 | Forecasting Returns over Multiple Time Horizons | 255 |
| 9.3.1 | Intrinsic Time | 255 |
| 9.3.2 | Model Structure | 256 |
| 9.3.3 | A Linear Combination of Nonlinear Indicators | 256 |
| 9.3.4 | Moving Averages, Momenta, and Indicators | 257 |
| 9.3.5 | Continuous Coefficient Update | 259 |
| 9.4 | Measuring Forecast Quality | 261 |
| 9.4.1 | Appropriate Measures of Forecast Accuracy | 262 |
| 9.4.2 | Empirical Results for the Multi-Horizon Model | 263 |
| 9.4.3 | Forecast Effectiveness in Intraday Horizons | 264 |

10

CORRELATION AND MULTIVARIATE RISK

| | | |
|---------------|----------------------------------------------------|-----|
| 10.1 | Introduction | 268 |
| 10.2 | Estimating the Dependence of Financial Time Series | 269 |
| 10.3 | Covolatility Weighting | 270 |
| 10.3.1 | Formulation of an Adjusted Correlation Measure | 272 |
| 10.3.2 | Monte Carlo and Empirical Tests | 274 |
| 10.4 | Stability of Return Correlations | 277 |
| 10.4.1 | Correlation Variations over Time | 278 |
| 10.4.2 | The Exponential Memory of Return Correlations | 282 |
| 10.5 | Correlation Behavior at High Data Frequencies | 287 |
| 10.6 | Conclusions | 293 |

11

TRADING MODELS

| | | |
|---------------|-------------------------------------------------------|-----|
| 11.1 | Introduction | 295 |
| 11.2 | Real-Time Trading Strategies | 297 |
| 11.2.1 | The Trading Model and Its Data-Processing Environment | 299 |
| 11.2.2 | Simulated Trader | 303 |
| 11.3 | Risk Sensitive Performance Measures | 304 |
| 11.3.1 | X_{eff} : A Symmetric Effective Returns Measure | 305 |
| 11.3.2 | R_{eff} : An Asymmetric Effective Returns Measure | 307 |
| 11.4 | Trading Model Algorithms | 309 |

LIST OF FIGURES

| | | |
|-------------|--------------------------------------------------|----|
| 1.1 | Size and data frequency of different samples | 4 |
| 1.2 | Models versus time scale | 5 |
| 1.3 | Volatility with daily versus high-frequency data | 7 |
| 3.1 | Types of time series operators | 36 |
| 3.2 | Interpolation methods | 38 |
| 3.3 | Overlapping time intervals | 48 |
| 3.4 | One week of USD-CHF prices | 53 |
| 3.5 | Moving average (MA) kernel | 61 |
| 3.6 | MA kernel on a logarithmic scale | 62 |
| 3.7 | Schematic differential kernel | 64 |
| 3.8 | Kernel of a differential operator | 65 |
| 3.9 | Decay of a differential kernel | 66 |
| 3.10 | Differential and return | 67 |
| 3.11 | Distribution of derivative operator | 68 |
| 3.12 | Annualized volatility as a moving norm | 71 |
| 3.13 | Moving moments of returns | 72 |
| 3.14 | Kernel of a windowed Fourier operator | 75 |
| 3.15 | Normed windowed Fourier transform | 77 |
| 3.16 | Microscopic volatility | 79 |
| 3.17 | Tick activity | 80 |
| 4.1 | Flowchart of a data-cleaning filter | 87 |

| | | |
|-------------|----------------------------------------------------------------------------------------------|-----|
| 4.2 | Schematic scalar filtering window | 103 |
| 5.1 | Short-term autocorrelation of returns | 123 |
| 5.2 | Comparison between quoted and transaction spreads | 125 |
| 5.3 | Scaling law exponent as a function of time | 128 |
| 5.4 | Seasonality in the interest rates | 129 |
| 5.5 | Intraday distribution of 15-min mean changes of absolute returns | 130 |
| 5.6 | Cumulative distributions of 10-min, 1-day, and 1-week USD-JPY returns | 136 |
| 5.7 | Order statistics for Student- t distribution | 139 |
| 5.8 | Scaling law for USD-JPY and GBP-USD | 149 |
| 5.9 | Wavelet variance at different scales | 160 |
| 5.10 | Autocorrelations of hourly returns, absolute returns, and squared returns | 161 |
| 5.11 | Autocorrelation as a function of the power of the absolute returns | 162 |
| 5.12 | Hourly intraday and intraweek distribution of absolute return, spread and the tick frequency | 165 |
| 5.13 | Intraday analysis of Eurofutures | 168 |
| 5.14 | Deterministic volatility of Eurofutures | 169 |
| 5.15 | Cumulative distributions of spreads | 172 |
| 6.1 | The USD-DEM intraweek activity pattern | 178 |
| 6.2 | Activity functions of geographical market components | 182 |
| 6.3 | Histograms of the average hourly activity for USD-JPY and USD-CHF | 185 |
| 6.4 | The activity model for USD-JPY and USD-CHF | 186 |
| 6.5 | Comparison of tick activity and volatility for different data sources | 188 |
| 6.6 | The ϑ -time versus physical time for USD-DEM | 190 |
| 6.7 | Hourly returns of USD-DEM in physical and ϑ -time | 191 |
| 6.8 | Seasonality and wavelet filtering | 194 |
| 6.9 | Autocorrelations of the 5-min absolute returns for USD-DEM and USD-JPY | 195 |
| 6.10 | Autocorrelations of the 5-min filtered absolute returns for USD-DEM and USD-JPY | 195 |
| 7.1 | Interaction of trader groups | 199 |
| 7.2 | The bias of realized volatility | 201 |
| 7.3 | The residual bias of bias-corrected realized volatility | 203 |
| 7.4 | Autocorrelation function of USD-DEM in physical-time | 205 |
| 7.5 | Autocorrelation function of USD-DEM in ϑ -time | 206 |
| 7.6 | USD-DEM autocorrelations from daily data | 208 |
| 7.7 | Coarse and fine volatilities | 212 |
| 7.8 | Asymmetric lagged correlation for USD-DEM | 214 |

| | | |
|--------------|----------------------------------------------------------------------|-----|
| 7.9 | Asymmetric lagged correlation for Euromark IR futures | 216 |
| 7.10 | Asymmetry of lagged correlation | 217 |
| 7.11 | Conditional autocorrelation of returns | 218 |
| 8.1 | Estimated and theoretical GARCH coefficients in business time | 226 |
| 8.2 | Estimated and theoretical GARCH coefficients in ϑ -time | 227 |
| 8.3 | GARCH estimates on a moving sample | 230 |
| 8.4 | Moment conditions for a HARCH(2) process | 234 |
| 8.5 | Impacts of market components for HARCH processes | 241 |
| 9.1 | Standard RiskMetrics volatility at different daytimes | 252 |
| 9.2 | Momentum indicator for forecasting | 258 |
| 10.1 | Autocorrelation of absolute returns for USD-DEM | 277 |
| 10.2 | Linear correlation coefficients for USD/DEM/NLG | 280 |
| 10.3 | Linear correlation coefficients for USD/DEM/GBP | 281 |
| 10.4 | Linear correlation coefficients for USD/DEM/ITL | 282 |
| 10.5 | Linear correlation coefficients for DJIA/AMEX | 283 |
| 10.6 | Linear correlation coefficients for USD 3-6M/DEM 3-6M | 284 |
| 10.7 | Linear correlation coefficients for DEM 3-6M/DEM 9-12M | 285 |
| 10.8 | Autocorrelations of correlation coefficients | 286 |
| 10.9 | Exponential decay of the autocorrelation of correlation coefficients | 287 |
| 10.10 | Correlation coefficients as a function of return time interval | 289 |
| 10.11 | Correlation versus logarithmic return time interval | 290 |
| 10.12 | Correlation stabilization intervals versus data frequencies | 292 |
| 11.1 | Data flow within a real-time trading model | 298 |
| 11.2 | Crossover operator | 313 |
| 11.3 | Syntactic restrictions for basic arithmetic operators | 314 |
| 11.4 | Total return of a portfolio of 10 O&A trading models | 341 |
| 11.5 | Set of feasible portfolios with currency hedging | 342 |

LIST OF TABLES

| | | |
|-------------|----------------------------------------------------------------------|-----|
| 2.1 | The traditional FAFX page of Reuters | 16 |
| 2.2 | FX data frequency | 18 |
| 4.1 | Data cleaning filter structure | 89 |
| 4.2 | Credibility addition | 90 |
| 4.3 | Trust capital as a function of price move size and time interval | 96 |
| 4.4 | Active periods of the three generic markets | 101 |
| 4.5 | Data cleaning filter parameters | 117 |
| 4.6 | Data cleaning rejection rates | 119 |
| 5.1 | Moments of return distributions for FX rates | 133 |
| 5.2 | Moments of return distributions for FX cross rates | 134 |
| 5.3 | Tail index of FX returns | 140 |
| 5.4 | Tail index of FX cross rates | 141 |
| 5.5 | Tail index of spot interest rates | 143 |
| 5.6 | Estimated tail index for different data frequencies and sample sizes | 144 |
| 5.7 | Extreme risks in the FX market | 146 |
| 5.8 | Drift exponents for FX rates | 150 |
| 5.9 | Drift exponents for Eurofutures | 151 |
| 5.10 | Timezone conversion table | 163 |
| 5.11 | Average number of ticks versus day of the week | 164 |
| 5.12 | Average volatility versus day of the week | 166 |

| | | |
|--------------|-------------------------------------------------------------------------|-----|
| 5.13 | Correlation coefficients between activity measures | 167 |
| 5.14 | Average spreads versus day of the week | 171 |
| 6.1 | Definition of the three generic markets | 180 |
| 6.2 | The ϑ -time parameter estimates for the three generic markets | 184 |
| 6.3 | The volatility ratio for the quality of the ϑ -scale | 192 |
| 7.1 | Difference between lagged correlations | 213 |
| 8.1 | Results of a GARCH(1,1) estimation in business time | 228 |
| 8.2 | Results of a GARCH(1,1) estimation in ϑ -time | 229 |
| 8.3 | Market components of a HARCH process | 236 |
| 8.4 | HARCH coefficients for USD-DEM | 240 |
| 8.5 | Results of the EMA-HARCH for the LIFFE Three-Month Euromark | 242 |
| 8.6 | Volatility forecasting performance for USD-DEM | 246 |
| 9.1 | The sampling periods of the forecast study | 264 |
| 9.2 | Forecast quality for 10 FX rates against the USD | 265 |
| 9.3 | Forecast quality for 10 FX cross rates | 266 |
| 9.4 | Significance of the forecast quality for 20 FX rates | 267 |
| 10.1 | Correlations from Monte Carlo simulations | 275 |
| 10.2 | Data sampling for correlation as function of time | 278 |
| 10.3 | Mean values, variances, maxima and minima of correlation | 279 |
| 10.4 | Estimation results of the autocorrelation of correlation | 288 |
| 10.5 | Correlation results characterizing the Epps effect | 291 |
| 11.1 | Market time constraints | 299 |
| 11.2 | Trading model results versus tree complexity | 316 |
| 11.3 | Performance comparison between models | 324 |
| 11.4 | Performance comparison between markets | 325 |
| 11.5 | The best X_{eff} as a function of opening hours | 326 |
| 11.6 | p -value Comparisons | 331 |
| 11.7 | Random walk Simulations for USD-DEM | 332 |
| 11.8 | GARCH(1,1) parameter estimates | 334 |
| 11.9 | GARCH(1,1) simulations for USD-DEM | 335 |
| 11.10 | AR(4)-GARCH(1,1) parameter estimates | 337 |
| 11.11 | AR(4)-GARCH(1,1) simulations for USD-DEM | 338 |
| 11.12 | Portfolio performance of O&A trading models | 340 |

PREFACE

This book presents a unified view of high-frequency time series methods with a particular emphasis on foreign exchange markets as well as interest rate spot and futures markets. The scope of this book is also applicable to other markets, such as equity and commodity markets.

As the archetype of financial markets, the foreign exchange market is the largest financial market worldwide. It involves dealers in different geographic locations, time zones, and working hours who have different time horizons, home currencies, information access, transaction costs, and other institutional constraints. The time horizons vary from intraday dealers, who close their positions every evening, to long-term investors and central banks. In this highly complex and heterogeneous market structure, the market participants are faced with different constraints and use different strategies to reach their financial goals, such as by maximizing their profits or maximizing their utility function after adjusting for market risk.

This book provides a framework to the analysis, modeling, and inference of high-frequency financial time series. It begins with the elementary foundations and definitions needed for studying the fundamental properties of high-frequency financial time series. It extends into the adaptive data-cleaning issues, treatment of seasonal volatility, and modeling of intraday volatility. Fractal properties of the high-frequency financial time series are found and explored, and an intrinsic time is used to construct forecasting models. The book provides a detailed study of how the adopted framework can be effectively utilized to build econometric models of

the price-formation process. Going beyond the price-formation process, the book presents the techniques to construct real-time trading models for financial assets.

It is designed for those who might be starting research in the area as well as for those who are interested in appreciating the statistical and econometric theory that underlies high-frequency financial time series modeling. The targeted audience includes finance professionals, including risk managers and research professionals in the public and private sectors; those taking graduate courses in finance, economics, econometrics, statistics, and time series analysis; and advanced MBA students. Because the high-frequency finance field is relatively new and the literature is scattered in a wide range of academic and nonacademic platforms, this book aims to provide a uniform treatment of the field and an easily accessible platform to high-frequency financial time series analysis — an exciting new field of research.

With the development of this field, a huge new area of research has been initiated, where work has hardly started. This work could not be more fascinating, and a number of discoveries are waiting to be made. We expect research to increase in this field, as people start to understand how these insights can dramatically improve risk-adjusted performances in asset management, market making, and treasury functions and be the foundation for other applications, such as an early warning system of financial markets.

Michel M. Dacorogna
Ramazan Gençay
Ulrich A. Müller
Richard B. Olsen
Olivier V. Pictet

ACKNOWLEDGMENTS

We should start by acknowledging that 15 years ago, when our research team at Olsen & Associates (O&A) first began using the amazing magnifying glass provided by high-frequency data to see if we could uncover possible patterns in the financial markets, none of us anticipated just how expansive the effort would become.

With the publication of this book originating from our work, sincere thanks are due to so many that we can only hope we have recognized most of the colleagues and friends who have advanced our work. Their help, their encouragement, their criticism, and their friendship have contributed to the style of teamwork we always favored.

We begin with Matthias Schwarz, a biology student who computed the first scaling law working with us in the autumn of 1986. Our first academic visitor was Claude Morgenegg, coming from the University of Geneva, who taught our group of physicists the right language to use to reach the economists. Thanks also to Casper de Vries, who opened up for us the world of extreme value theory; Cindy L. Gauveau, who prepared forecasting models for foreign exchange rates and also brought the economic touch to our work; Rakhal Davé, for his explorations of the LeBaron effect; Marco Tomassini and Bastien Chopard, who brought to our attention the genetic algorithms; Mark Lundin and his correlation studies; Gennady Samorodnitsky and Paul Embrechts, who were able to prove the sufficiency of the stationarity condition of HARCH processes; Giuseppe Ballocci, who led us into the research on interest rate futures; and Wolfgang

Breymann, who has extended the ϑ -time concept and developed the idea of a heterogeneous market in his cascade model. A particular thanks goes to Gilles Zumbach, who has contributed many graphs to this book. He continues the work, bringing it to new levels and uncovering many more properties with the powerful operator framework and the software tools in C++ that he has developed over the years.

We also want to thank our colleagues who joined us at a later stage and have already reached out for other adventures: Lars Jaeger, Thomas Domenig, Peter Rice, and Hoss Hauksson. Our thanks extend also to Jørgen Olsen, whose wisdom and vast scientific culture has enlightened our seminars and whose road map for building a Richter scale for financial markets we implemented.

One very important and enriching experience has been the visits of many students who spent time with us and brought along their enthusiasm and eagerness to learn: Dominique Guillaume; Lukas Pulver; Petra Korndorfer; Markus P. Herrchen; Jens Richelsen; Christian Jost; Jürg S. Füssler; Retus G. Sgier; Alexander Dimai; Jonathan Dawes; Jakob E. von Weizsäcker; Philipp Hartmann; Cătălin Stărică; Barbara Piccinato; Carl Hopman; Peter Rice, who later joined our research team; Simone Deparis; Fulvio Corsi; and Paul Lynch. Without them, we would never have been able to explore so many different time series and to accomplish so many studies.

Many of our academic friends around the world visited us and understood early on the interest in research on this type of data. They provided us with encouragement to continue and the sense that we were working in the right direction. Hermann Garbers was the first to invite us to give a seminar at the University of Zurich, where we presented the scaling law in December 1988. From Benoît Mandelbrot in 1989 to Gennady Samorodnitsky just prior to the publication of this book, we have been fortunate to share time and work at O&A with some fine scientists: Tim Bollerslev, William Brock, Hans Bühlmann, Peter Bühlmann, Frank K. Diebold, Christian Dunis, Rüdiger Frey, Hélyette Geman, Charles Goodhart, Rudolf Kalman, Hans Rudolf Lerche, Bruce Mizrach, John Moody, Salih Neftçi, Wolfgang Polasek, Remo Schnidrig, Albert N. Shiryaev, Gerhard Stahl, Massimo Tivegna, Murad Taqqu, Walter Wasserfallen, Andreas Weigend, and Diethelm Würtz. We would like to thank especially Charles Goodhart, whose support and insights led to the O&A “High-Frequency Data in Finance” conferences, but also Richard Baillie, Tim Bollerslev, Rob Engle, Joel Hasbrouck, Michael Melvin, and Maureen O’Hara. Gerhard Stahl has been a great partner in exploring new issues of risk management. His scientific rigor was always refreshing in this field where ad hoc arguments often dominate. Michel Dacorogna would like to thank particularly Blake LeBaron for the many e-mail exchanges we have had over the years on our research. They were always stimulating and encouraged us to think deeper into the problems and find the connections with the traditional economic approach. Manfred Härter and Mico Loretan have been always supportive of our work and have brought to us many ideas and opportunities for presenting it. Ramo Gençay would like to thank William Brock, Dee Dechert, Murray Frank,

Blake LeBaron, and Thanasis Stengos for many exciting research conversations, and Michael Charette, Ron Meng and Tibor Toronyi for research support. Ulrich Müller would like to thank Günter Schwarz for his contribution to our understanding of portfolio theory.

It is also clear that without the help and the dedication of our software team we would not have been able to access a database of such quantity and quality, covering more than 14 years of tick-by-tick prices. From Rob J. Nagler to Kris Meissner through J. Robert Ward, William H. Kelly, Daniel P. Smith, Martin Lichtin, Devon Bowen, and Michael Stumm, we learned the subtleties of object-oriented programming and have enjoyed their constant support in our efforts to make sense of all that we were seeing. Paul Breslaw has been so helpful with the data and improving our English. Our thanks go also to our friends from the operation group who kept alive our system, especially Jorge Mota, Jeff Courtade, and Gary Swofford. Whenever we had a problem with bulbs going out of function or air conditioning not working (especially when it was needed most), Filippo Guglielmo would always be here to solve it.

The trading model developments would not have been so interesting, nor so close to reality, without the contribution of our help desk: Pius Dall'Acqua, Stephan Schlatter, and last, but not least, Bernard Hechinger and his deep knowledge of the microstructure of financial markets. His interest in our models has brought us to rethink many aspects of their strategies and implement some of his ideas. In terms of market knowledge, we especially want to thank Dean LeBaron, whose vast experience and enthusiasm for new developments is a source of inspiration and encouragement. Our customers also brought many ideas to us, especially in the group who participated in the development of the interest rate project: Michael Brockmann, Dieter Heitkamp, Luciano Steve, and Giuseppe Ciliberto. We always enjoyed the exchanges with the practitioners who understood the need for a scientific approach to the markets. Another example of this fertile interaction was with Monique Donders, Tjark Tjin, and Marcel Vernooy during our project of building a currency overlay product based on trading models. Special thanks go also to Daniel Huber, who opened up so many doors for us.

There is no question that we benefited greatly from the structure and organization provided by our different administrative assistants over the years. Karin Jost, Rosemarie Arnold-Becker, and Melanie Käslin all brought a strong sense of service and dedication that made our teamwork possible.

The process of writing a book with many authors is complex and demanding but also very rewarding because it gave us the occasion to discuss, deepen our understanding of the matters and interact with interesting people. In this process, Scott Bentley's (the senior editor of Academic Press) help and feedback have been important for keeping the level of motivation high and for the success of this project. The care of Amy Hendrickson for many \LaTeX formatting problems of a book of more than 400 pages and containing so many figures was essential for the resulting appearance of this book.

Before closing this page of gratitude, we do not want to forget Dina Weidmann and Elisa Guglielmo, who cooked so many fine dishes with the Italian touch and make O&A's famous "Friday family lunches" a genuine gourmet experience. Faced with mountains of data to unravel, this lovely tradition warmed the soul. Grazie.

Michel M. Dacorogna
Ramazan Gençay
Ulrich A. Müller
Richard B. Olsen
Olivier V. Pictet

INTRODUCTION

I.1 MARKETS: THE SOURCE OF HIGH-FREQUENCY DATA

A famous climber, when asked why he was willing to put his life in danger to climb dangerous summits, answered: “Because they are there.” We would be tempted to give the same answer when people ask us why we take so much pain in dealing with high-frequency data. The reason is simple: financial markets are the source of high-frequency data. The original form of market prices is tick-by-tick data: each “tick” is one logical unit of information, like a quote or a transaction price (see Section 2.1). By nature these data are irregularly spaced in time. Liquid markets generate hundreds or thousands of ticks per business day. Data vendors like Reuters transmit more than 275,000 prices per day for foreign exchange spot rates alone.

Thus high-frequency data should be the primary object of research for those who are interested in understanding financial markets. Especially so, because practitioners determine their trading decisions by observing high-frequency or tick-by-tick data. Yet most of the studies published in the financial literature deal with low-frequency, regularly spaced data. There are two main reasons for this. First, it is still rather costly and time-consuming to collect, collate, store, retrieve, and manipulate high-frequency data. That is why most of the available

financial data are at daily or lower frequency. The second reason is somehow more subtle but still quite important: most of the statistical apparatus has been developed and thought for homogeneous (i.e., equally spaced in time) time series. There is little work done to adapt the methods to data that arrive at random time intervals. Unfortunately in finance, regularly spaced data are not original data but artifacts derived from the original market prices. Nowadays with the development of computer technology, data availability is becoming less and less of a problem. For instance, most of the exchanges and especially those that trade electronically would gladly provide tick-by-tick data to interested parties. Data vendors have themselves improved their data structures and provide their users with tools to collect data for over-the-counter (OTC) markets. Slowly, high-frequency data are becoming a fantastic experimental bench for understanding market microstructure and more generally for analyzing financial markets.

That leaves the researcher with the problems of dealing with such vast amounts of data using the right mathematical tools and models. This is precisely the subject of this book.

1.2 METHODOLOGY OF HIGH-FREQUENCY RESEARCH

From the beginning, our approach has been to apply the experimental method which has been highly successful in “hard” sciences.¹ It consists of three steps, the first one being to explore the data in order to discover the fundamental statistical properties they exhibit with a minimum set of assumptions. This is often called finding the “stylized facts” in the econometric or finance literature. This first step was in fact not so important in the economic literature, because the sparseness of data made it either relatively simple or uninteresting due to the statistical uncertainty.

The second step is to use all of these empirical facts to formulate adequate models. By adequate models, we do not mean models that come from hand-waving arguments about the markets, but rather models that are directly inspired by the empirical regularities encountered in the data. It is the point where our understanding of market behavior and reality of the data properties should meet. There have been many debates between the time series approach and microstructure approach. The first one relying more on modeling the statistical properties of the data and the latter concentrating on modeling market behavior. Both approaches have their value and high-frequency data might be able to reconcile them by enabling us to actually test the microstructure models, Hasbrouck (1998); Rydberg and Shephard (1998).

The third step, of course, is to verify whether these models satisfactorily reproduce the stylized facts found in the data. The ultimate goal is not only a good descriptive model but the ability to produce reasonable *predictions* of future movements or risks and to integrate these tools into practical applications, such

¹ We refer here to experimental sciences such as physics, chemistry, or biology.

as risk management tools or option pricing algorithms. For decades, practitioners have been developing so-called technical analysis, which is a kind of empirical time series analysis based on rudimentary analytical tools. Although some new academic research has analyzed these trading rules,² they remain controversial and are looked down upon. We hope that this book will put on a new footing many ideas that have been developed in technical analysis.

We have organized this book along the same lines, we first present the empirical regularities, then we construct models, and lastly we test their power to predict market outcomes.

The novelty of high-frequency data demands to take such an approach. This was not usual in econometrics because so little data were available until the late 1980s. It was quite natural that the researcher's emphasis was to make sure that the methodology was correct in order to obtain the most information out of the sparse data that were available. Only recently the research community in this field has recognized the importance of the first step: finding empirical facts. This step can already be good research in its own right. A good example is the recent paper by Andersen *et al.* (2001), where the authors explore in detail the distributional properties of volatility computed from high-frequency data.

Thanks to the development of electronic trading and the existence of various data providers also on the Internet, it is now possible to follow the price formation in real-time. Ideally, the analysis and modeling of the price-generation process should, in real-time, produce results that add value to the raw data. There is strong demand from the market to have, next to the current price, a good assessment of the current risk of the financial asset as well as a reasonable prediction of its future movement. This means that the models should be made amenable to real-time computations and updates. Techniques for doing so will be presented in the remainder of the book. It is possible to develop methods that allow for the easy computation of models and can thus provide almost instantaneous reaction to market events. Although quite popular among practitioners who want to analyze the past developments of prices, those techniques have had little echo, until now, in the academic world. Very few research papers have studied the statistical foundations and properties of those "technical indicators." In this book (Chapter 3) we provide a unified platform for these methods.

1.3 DATA FREQUENCY AND MARKET INFORMATION

Relating the type of data available for researchers, the effects and the models that are discovered and developed with these different samples, provides insight into the development of research in finance. Figure 1.1 illustrates the sample size versus the measurement frequency of some well-known data sets used in finance. The

² Among others, here is a list of interesting papers on the issue of technical trading models: Neftci (1991), Brock *et al.* (1992), Taylor and Allen (1992), Levich and Thomas (1993b), Gençay and Stengos (1998), Gençay (1998a,b), Frances and van Griensven (1998), Allen and Karjalainen (1999), Gençay (1999), LeBaron (1999a), Sullivan *et al.* (1999), and Gençay *et al.* (2001c, 2002).




FIGURE 1.1 Available data samples with their typical sizes and frequency. The sample size and the frequency are plotted on a logarithmic scale. The first point corresponds to the O&A database, the last one to the 700 years of yearly data analyzed by Froot *et al.* (1995), the second to its left to the cotton price data of Mandelbrot (1963), and the daily data are computed from the sample used in Ding *et al.* (1993) to show long memory in the S&P 500. The text refers to the effects discovered and analyzed in the different segments of these samples.

double logarithmic scale makes the points lie almost on a straight line. The data sample with the lowest frequency is the one used by Froot *et al.* (1995) of 700 years of annual commodity price data from England and Holland. Beyond 700 years, one is unlikely to find reliable economic or financial data.³ The data with the highest frequency is the Olsen & Associates (O&A) dataset of more than 14 years of high-frequency foreign exchange data. The tick-by-tick data are the highest frequency available. Between those two extremes, one finds the daily series of the Standard & Poors 500 from 1928 to 1991 used by Ding *et al.* (1993) or the monthly cotton prices used by Mandelbrot (1963) from 1880 to 1940. On this graph, we superimpose those effects that have been identified at these different time scales. One of the questions with data collected over very long periods is whether they really refer to the same phenomenon. Stock indices, for example, change their composition through time due to mergers or the demise of companies. When analyzing the price history of stock indices, the impact of these changes in

³ Data can be found in natural sciences such as weather data up to a few hundred thousand years.

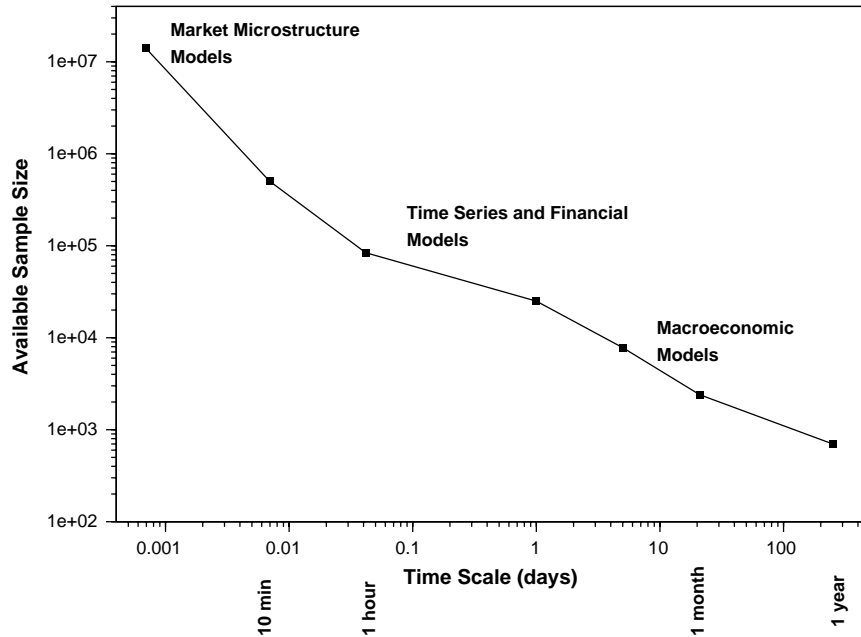


FIGURE 1.2 Available data samples with their typical sizes and frequency. The sample size and the frequency are plotted on a logarithmic scale. The text refers to the models developed and tested in the different segments of these samples.

composition is not obvious. We call this phenomenon the “breakdown of the permanence hypothesis.” It is difficult to assess the quality of any inference as the underlying process is not stationary over decades or centuries. At the other end of the frequency spectrum (i.e. with high-frequency data), we are confronted with the details of the price generation process, where other effects, such as how the data are transmitted and recorded in the data-base (see Chapter 4) have an impact. With data at frequencies of the order of one hour, a new problem arises, due to the fact that the earth turns and the impact of time zones, where the seasonality of volatility becomes very important (as we shall see in Chapter 5) and overshadows all other effects.

Figure 1.2 relates the data to the models that are typically developed and tested with them. The high-frequency data have opened great possibilities to test market microstructure models, while traditionally low-frequency data are used for testing macroeconomic models. In between lies the whole area of financial and time series modeling, which is typically studied with daily or monthly data as, for instance, option pricing or GARCH models. It is clear from this figure that we have a continuum of both samples and models. The antagonism that is sometimes encountered between time series and market microstructure approaches should slowly vanish with more and more studies combining both with high-frequency

data. Yet the challenge is still open to build models that are simple to implement and describe to a reasonable degree the empirical behavior of the data at all time scales.

1.4 NEW LEVELS OF SIGNIFICANCE

High-frequency data means a very large amount of data. The number of observations in one single day of a liquid market is equivalent to the number of daily data within 30 years. Statistically, the higher the number of independently measured observations, the higher is the degrees of freedom, which implies more precise estimators. The large amount of data allows us to distinguish between different models (model validation) with a higher statistical precision. New statistical methods become possible, for example, tail statistics to examine the probability of extreme events. Almost by definition, extreme events are rare and doing statistics on such extreme events is a challenge. With high-frequency data one can have samples with as many as 400,000 independent observations⁴ to study the 0.25% percentile and still have 1,000 observations with which to work. We shall see how important this is when we present the estimation of tail indices for return distributions. Similarly, when different models have to be ranked, the availability of a few hundred thousand observations allows us to find beyond a doubt which model provides the best description of the data-generating process (Müller *et al.*, 1997a).

Figure 1.3 demonstrates the importance of high-frequency data in model selection and inference within the context of Value-at-Risk (VaR) calculations. We report three different calculations all of which use the J. P. Morgan (1996) volatility model, which is in fact a 1-day volatility forecast as further discussed in Section 9.2. The three calculations differ in terms of the sampling and the data frequency. The Japanese volatility calculations are based on prices observed *daily* at 7 a.m. GMT, which corresponds to the afternoon Japanese time. The U.K. volatility calculations are based on prices measured *daily* at 5 p.m. GMT, which is the afternoon in the U. K. The high-frequency volatility calculations are based on the high-frequency tick-by-tick data recorded continuously on a 24-hour cycle. The top panel in Figure 1.3 reports the annualized volatility calculations and the bottom panel shows the underlying prices for January and February 1999. The top panel demonstrates that volatility can be extremely different depending on the time of the day at which it is measured with daily data. If observations are picked randomly once a day, the underlying volatility can be as small as 15% or as large as 22% for a given day and for the same currency. In mid-January 1999, the U.S. Dollar - Japanese Yen (USD-JPY) investors in the U.K. are assumed to be facing the risk of losing 56,676,400 USD in a portfolio of a hundred million USD with a 1% probability. In Japan, this risk would be reduced to 38,643,000 USD for the same day and for the same currency, a difference of approximately 18,000,000 USD between the two geographical locations! The utilization of high frequency leads to more robust

⁴ This approximately corresponds to 10 years of returns measured over 10 minutes.




FIGURE 1.3 Top panel: Annualized USD-JPY volatility computed with daily prices observed at 7 a.m. GMT (afternoon Japan, circles), 5 p.m. GMT (afternoon U.K., diamonds) and with high-frequency data (solid line). The data period is from January 1999 to February 1999. Bottom panel: The USD-JPY high-frequency price series from January 1999 to February 1999.

annualized volatility estimations by minimizing the influence of the random noise in the market.

Another aspect of this is the choice of model. With few data, one tends to favor the simpler models because they contain few parameters and because tests like the likelihood ratio test would strongly penalize the increase of parameters. Of course, simplicity is a desirable feature of theoretical models, but one should not seek simplicity at the cost of missing important features of the data-generating process. Sometimes, it is useful to explore more complicated (nonlinear) models, which may contain more parameters. This increasing complexity is strongly penalized when explored with low-frequency data because of the loss of degrees of freedom. In the case of high-frequency data, however, the penalty is relatively small because the abundance of the independently measured observations approximates an asymptotic environment.

Researchers who want to use many observations with low-frequency data are using, for instance, daily observations of the Dow Jones Industrials from January 1897 like Ding *et al.* (1993) or LeBaron (1999a). In such a case, one is entitled to ask if the authors are actually analyzing the same market over the years. The huge technological changes that we experienced during this century have certainly affected the New York Stock Exchange and one is never sure, how this and any reconfiguration of the index has affected the results. To the contrary, high-frequency studies can be done for limited sampling periods with reasonably large samples. The market properties within such periods are nearly unchanged. The results are less affected by structural breaks or shifts in the overall economy than low-frequency studies with samples of many years. This is clearly an advantage when determining microstructure effects but also when examining the stability of some properties over time.

1.5 INTERRELATING DIFFERENT TIME SCALES

High-frequency data open the way for studying financial markets at very different time scales, from minutes to years. This represents an aggregation factor of four to five orders of magnitude.⁵ Some empirical properties are similar at different scales, leading to fractal behaviors. Stylized facts observed for daily or weekly data gain additional weight when also observed with high significance for intraday data. An example of this is the long memory effect in 20-min absolute returns studied by Dacorogna *et al.* (1993). At the time, similar hyperbolic decay of the autocorrelation function was observed on daily returns in Ding *et al.* (1993). It is very difficult to distinguish rigorously in the data between long memory effects and regime shifts. Many mathematicians are working precisely on this problem such as Mansfield *et al.* (1999) and Mikosch and Starica (1999). Yet the fact that hyperbolic decay is empirically found at time scales that differ by two orders of magnitude in aggregation is definitely a sign that the process must include some long range dependence or that there are regime shifts at *all time scales*, which is equivalent.

Scaling properties and scaling laws have been new objects of study since the early work of Mandelbrot (1963) on cotton prices. In 1990, the research group of O&A published empirical studies of scaling properties extending from a few minutes to a few years (Müller *et al.*, 1990). These properties have shown remarkable stability over time (Guillaume *et al.*, 1997) and were found in other financial instruments like interest rates (Piccinato *et al.*, 1997). Mantegna and Stanley (1995) also found scaling behavior in the stock indices examined at high frequency. In a set of recent papers, Mandelbrot *et al.* (1997), Fisher *et al.* (1997) and Calvet *et al.* (1997) have derived a multifractal model based on the empirical scaling laws of different moments of the return distributions. Works on the scaling law of return

⁵ By order of magnitude we mean the number of times the time horizon must be multiplied by 10 to achieve the lower frequency. For instance, a weekly frequency is aggregated three orders of magnitude from 10 minutes data (one week is 1008 times 10 minutes).

volatility have been flourishing in the past few years often coming from physicists who started venturing in the field of finance calling themselves “econophysicists.” It is a sign that the field is moving toward a better understanding of aggregation properties. Unfortunately, the mathematical theory behind these empirical studies is not yet completely mature and there is still controversy regarding the significance of the scaling properties (LeBaron, 1999a; Bouchaud *et al.*, 2000). Thanks to high-frequency data, this kind of debate can now take place. The challenge is to develop models that simultaneously characterize the short-term *and* the long-term behaviors of a time series.