



FINANCIAL SIGNAL PROCESSING AND MACHINE LEARNING

EDITED BY

ALI N. AKANSU

SANJEEV R. KULKARNI

DMITRY M. MALIOUTOV



IEEE PRESS

WILEY

FINANCIAL SIGNAL PROCESSING AND MACHINE LEARNING

Edited by

Ali N. Akansu

New Jersey Institute of Technology, USA

Sanjeev R. Kulkarni

Princeton University, USA

Dmitry Malioutov

IBM T.J. Watson Research Center, USA



IEEE PRESS

WILEY

This edition first published 2016
© 2016 John Wiley & Sons, Ltd
First Edition published in 2016

Registered office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. It is sold on the understanding that the publisher is not engaged in rendering professional services and neither the publisher nor the author shall be liable for damages arising herefrom. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data applied for

ISBN: 9781118745670

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

Set in 10/12pt, TimesLTStd by SPi Global, Chennai, India.

Printed and bound in Singapore by Markono Print Media Pte Ltd

List of Contributors

Ali N. Akansu, New Jersey Institute of Technology, USA

Marco Cuturi, Kyoto University, Japan

Alexandre d'Aspremont, CNRS - Ecole Normale supérieure, France

Christine De Mol, Université Libre de Bruxelles, Belgium

Jianqing Fan, Princeton University, USA

Jun-ya Gotoh, Chuo University, Japan

Nicholas A. James, Cornell University, USA

Prabhanjan Kambadur, Bloomberg L.P., USA

Alexander Kreinin, Risk Analytics, IBM, Canada

Sanjeev R. Kulkarni, Princeton University, USA

Yuan Liao, University of Maryland, USA

Han Liu, Princeton University, USA

Matthew Lorig, University of Washington, USA

Aurélié C. Lozano, IBM T.J. Watson Research Center, USA

Ronny Luss, IBM T.J. Watson Research Center, USA

Dmitry Malioutov, IBM T.J. Watson Research Center, USA

David S. Matteson, Cornell University, USA

William B. Nicholson, Cornell University, USA

Ronnie Sircar, Princeton University, USA

Akiko Takeda, The University of Tokyo, Japan

Mustafa U. Torun, New Jersey Institute of Technology, USA

Stan Uryasev, University of Florida, USA

Onur Yilmaz, New Jersey Institute of Technology, USA

Preface

This edited volume collects and unifies a number of recent advances in the signal-processing and machine-learning literature with significant applications in financial risk and portfolio management. The topics in the volume include characterizing statistical dependence and correlation in high dimensions, constructing effective and robust risk measures, and using these notions of risk in portfolio optimization and rebalancing through the lens of convex optimization. It also presents signal-processing approaches to model return, momentum, and mean reversion, including both theoretical and implementation aspects. Modern finance has become global and highly interconnected. Hence, these topics are of great importance in portfolio management and trading, where the financial industry is forced to deal with large and diverse portfolios in a variety of asset classes. The investment universe now includes tens of thousands of international equities and corporate bonds, and a wide variety of other interest rate and derivative products—often with limited, sparse, and noisy market data.

Using traditional risk measures and return forecasting (such as historical sample covariance and sample means in Markowitz theory) in high-dimensional settings is fraught with peril for portfolio optimization, as widely recognized by practitioners. Tools from high-dimensional statistics, such as factor models, eigen-analysis, and various forms of regularization that are widely used in real-time risk measurement of massive portfolios and for designing a variety of trading strategies including statistical arbitrage, are highlighted in the book. The dramatic improvements in computational power and special-purpose hardware such as field programmable gate arrays (FPGAs) and graphics processing units (GPUs) along with low-latency data communications facilitate the realization of these sophisticated financial algorithms that not long ago were “hard to implement.”

The book covers a number of topics that have been popular recently in machine learning and signal processing to solve problems with large portfolios. In particular, the connections between the portfolio theory and sparse learning and compressed sensing, robust optimization, non-Gaussian data-driven risk measures, graphical models, causal analysis through temporal-causal modeling, and large-scale copula-based approaches are highlighted in the book.

Although some of these techniques already have been used in finance and reported in journals and conferences of different disciplines, this book attempts to give a unified treatment from a common mathematical perspective of high-dimensional statistics and convex optimization. Traditionally, the academic quantitative finance community did not have much overlap with the signal and information-processing communities. However, the fields are seeing more interaction, and this trend is accelerating due to the paradigm in the financial sector which has

embraced state-of-the-art, high-performance computing and signal-processing technologies. Thus, engineers play an important role in this financial ecosystem. The goal of this edited volume is to help to bridge the divide, and to highlight machine learning and signal processing as disciplines that may help drive innovations in quantitative finance and electronic trading, including high-frequency trading.

The reader is assumed to have graduate-level knowledge in linear algebra, probability, and statistics, and an appreciation for the key concepts in optimization. Each chapter provides a list of references for readers who would like to pursue the topic in more depth. The book, complemented with a primer in financial engineering, may serve as the main textbook for a graduate course in financial signal processing.

We would like to thank all the authors who contributed to this volume as well as all of the anonymous reviewers who provided valuable feedback on the chapters in this book. We also gratefully acknowledge the editors and staff at Wiley for their efforts in bringing this project to fruition.

Contents

List of Contributors	xiii
Preface	xv
1 Overview	1
<i>Ali N. Akansu, Sanjeev R. Kulkarni, and Dmitry Malioutov</i>	
1.1 Introduction	1
1.2 A Bird's-Eye View of Finance	2
1.2.1 Trading and Exchanges	4
1.2.2 Technical Themes in the Book	5
1.3 Overview of the Chapters	6
1.3.1 Chapter 2: "Sparse Markowitz Portfolios" by Christine De Mol	6
1.3.2 Chapter 3: "Mean-Reverting Portfolios: Tradeoffs between Sparsity and Volatility" by Marco Cuturi and Alexandre d'Aspremont	7
1.3.3 Chapter 4: "Temporal Causal Modeling" by Prabhanjan Kambadur, Aurélie C. Lozano, and Ronny Luss	7
1.3.4 Chapter 5: "Explicit Kernel and Sparsity of Eigen Subspace for the AR(1) Process" by Mustafa U. Torun, Onur Yilmaz and Ali N. Akansu	7
1.3.5 Chapter 6: "Approaches to High-Dimensional Covariance and Precision Matrix Estimation" by Jianqing Fan, Yuan Liao, and Han Liu	7
1.3.6 Chapter 7: "Stochastic Volatility: Modeling and Asymptotic Approaches to Option Pricing and Portfolio Selection" by Matthew Lorig and Ronnie Sircar	7
1.3.7 Chapter 8: "Statistical Measures of Dependence for Financial Data" by David S. Matteson, Nicholas A. James, and William B. Nicholson	8
1.3.8 Chapter 9: "Correlated Poisson Processes and Their Applications in Financial Modeling" by Alexander Kreinin	8
1.3.9 Chapter 10: "CVaR Minimizations in Support Vector Machines" by Junya Gotoh and Akiko Takeda	8
1.3.10 Chapter 11: "Regression Models in Risk Management" by Stan Uryasev	8
1.4 Other Topics in Financial Signal Processing and Machine Learning	9
References	9

2	Sparse Markowitz Portfolios	11
	<i>Christine De Mol</i>	
2.1	Markowitz Portfolios	11
2.2	Portfolio Optimization as an Inverse Problem: The Need for Regularization	13
2.3	Sparse Portfolios	15
2.4	Empirical Validation	17
2.5	Variations on the Theme	18
	2.5.1 <i>Portfolio Rebalancing</i>	18
	2.5.2 <i>Portfolio Replication or Index Tracking</i>	19
	2.5.3 <i>Other Penalties and Portfolio Norms</i>	19
2.6	Optimal Forecast Combination	20
	Acknowledgments	21
	References	21
3	Mean-Reverting Portfolios	23
	<i>Marco Cuturi and Alexandre d'Aspremont</i>	
3.1	Introduction	23
	3.1.1 <i>Synthetic Mean-Reverting Baskets</i>	24
	3.1.2 <i>Mean-Reverting Baskets with Sufficient Volatility and Sparsity</i>	24
3.2	Proxies for Mean Reversion	25
	3.2.1 <i>Related Work and Problem Setting</i>	25
	3.2.2 <i>Predictability</i>	26
	3.2.3 <i>Portmanteau Criterion</i>	27
	3.2.4 <i>Crossing Statistics</i>	28
3.3	Optimal Baskets	28
	3.3.1 <i>Minimizing Predictability</i>	29
	3.3.2 <i>Minimizing the Portmanteau Statistic</i>	29
	3.3.3 <i>Minimizing the Crossing Statistic</i>	29
3.4	Semidefinite Relaxations and Sparse Components	30
	3.4.1 <i>A Semidefinite Programming Approach to Basket Estimation</i>	30
	3.4.2 <i>Predictability</i>	30
	3.4.3 <i>Portmanteau</i>	31
	3.4.4 <i>Crossing Stats</i>	31
3.5	Numerical Experiments	32
	3.5.1 <i>Historical Data</i>	32
	3.5.2 <i>Mean-reverting Basket Estimators</i>	33
	3.5.3 <i>Jurek and Yang (2007) Trading Strategy</i>	33
	3.5.4 <i>Transaction Costs</i>	33
	3.5.5 <i>Experimental Setup</i>	36
	3.5.6 <i>Results</i>	36
3.6	Conclusion	39
	References	39

4	Temporal Causal Modeling	41
	<i>Prabhanjan Kambadur, Aurélie C. Lozano, and Ronny Luss</i>	
4.1	Introduction	41
4.2	TCM	46
4.2.1	<i>Granger Causality and Temporal Causal Modeling</i>	46
4.2.2	<i>Grouped Temporal Causal Modeling Method</i>	47
4.2.3	<i>Synthetic Experiments</i>	49
4.3	Causal Strength Modeling	51
4.4	Quantile TCM (Q-TCM)	52
4.4.1	<i>Modifying Group OMP for Quantile Loss</i>	52
4.4.2	<i>Experiments</i>	53
4.5	TCM with Regime Change Identification	55
4.5.1	<i>Model</i>	56
4.5.2	<i>Algorithm</i>	58
4.5.3	<i>Synthetic Experiments</i>	60
4.5.4	<i>Application: Analyzing Stock Returns</i>	62
4.6	Conclusions	63
	References	64
5	Explicit Kernel and Sparsity of Eigen Subspace for the AR(1) Process	67
	<i>Mustafa U. Torun, Onur Yilmaz, and Ali N. Akansu</i>	
5.1	Introduction	67
5.2	Mathematical Definitions	68
5.2.1	<i>Discrete AR(1) Stochastic Signal Model</i>	68
5.2.2	<i>Orthogonal Subspace</i>	69
5.3	Derivation of Explicit KLT Kernel for a Discrete AR(1) Process	72
5.3.1	<i>A Simple Method for Explicit Solution of a Transcendental Equation</i>	73
5.3.2	<i>Continuous Process with Exponential Autocorrelation</i>	74
5.3.3	<i>Eigenanalysis of a Discrete AR(1) Process</i>	76
5.3.4	<i>Fast Derivation of KLT Kernel for an AR(1) Process</i>	79
5.4	Sparsity of Eigen Subspace	82
5.4.1	<i>Overview of Sparsity Methods</i>	83
5.4.2	<i>pdf-Optimized Midtread Quantizer</i>	84
5.4.3	<i>Quantization of Eigen Subspace</i>	86
5.4.4	<i>pdf of Eigenvector</i>	87
5.4.5	<i>Sparse KLT Method</i>	89
5.4.6	<i>Sparsity Performance</i>	91
5.5	Conclusions	97
	References	97

6	Approaches to High-Dimensional Covariance and Precision Matrix Estimations	100
	<i>Jianqing Fan, Yuan Liao, and Han Liu</i>	
6.1	Introduction	100
6.2	Covariance Estimation via Factor Analysis	101
6.2.1	<i>Known Factors</i>	103
6.2.2	<i>Unknown Factors</i>	104
6.2.3	<i>Choosing the Threshold</i>	105
6.2.4	<i>Asymptotic Results</i>	105
6.2.5	<i>A Numerical Illustration</i>	107
6.3	Precision Matrix Estimation and Graphical Models	109
6.3.1	<i>Column-wise Precision Matrix Estimation</i>	110
6.3.2	<i>The Need for Tuning-insensitive Procedures</i>	111
6.3.3	<i>TIGER: A Tuning-insensitive Approach for Optimal Precision Matrix Estimation</i>	112
6.3.4	<i>Computation</i>	114
6.3.5	<i>Theoretical Properties of TIGER</i>	114
6.3.6	<i>Applications to Modeling Stock Returns</i>	115
6.3.7	<i>Applications to Genomic Network</i>	118
6.4	Financial Applications	119
6.4.1	<i>Estimating Risks of Large Portfolios</i>	119
6.4.2	<i>Large Panel Test of Factor Pricing Models</i>	121
6.5	Statistical Inference in Panel Data Models	126
6.5.1	<i>Efficient Estimation in Pure Factor Models</i>	126
6.5.2	<i>Panel Data Model with Interactive Effects</i>	127
6.5.3	<i>Numerical Illustrations</i>	130
6.6	Conclusions	131
	References	131
7	Stochastic Volatility	135
	<i>Matthew Lorig and Ronnie Sircar</i>	
7.1	Introduction	135
7.1.1	<i>Options and Implied Volatility</i>	136
7.1.2	<i>Volatility Modeling</i>	137
7.2	Asymptotic Regimes and Approximations	141
7.2.1	<i>Contract Asymptotics</i>	142
7.2.2	<i>Model Asymptotics</i>	142
7.2.3	<i>Implied Volatility Asymptotics</i>	143
7.2.4	<i>Tractable Models</i>	145
7.2.5	<i>Model Coefficient Polynomial Expansions</i>	146
7.2.6	<i>Small “Vol of Vol” Expansion</i>	152
7.2.7	<i>Separation of Timescales Approach</i>	152
7.2.8	<i>Comparison of the Expansion Schemes</i>	154
7.3	Merton Problem with Stochastic Volatility: Model Coefficient Polynomial Expansions	155

7.3.1	<i>Models and Dynamic Programming Equation</i>	155
7.3.2	<i>Asymptotic Approximation</i>	157
7.3.3	<i>Power Utility</i>	159
7.4	Conclusions	160
	Acknowledgements	160
	References	160
8	Statistical Measures of Dependence for Financial Data	162
	<i>David S. Matteson, Nicholas A. James, and William B. Nicholson</i>	
8.1	Introduction	162
8.2	Robust Measures of Correlation and Autocorrelation	164
8.2.1	<i>Transformations and Rank-Based Methods</i>	166
8.2.2	<i>Inference</i>	169
8.2.3	<i>Misspecification Testing</i>	171
8.3	Multivariate Extensions	174
8.3.1	<i>Multivariate Volatility</i>	175
8.3.2	<i>Multivariate Misspecification Testing</i>	176
8.3.3	<i>Granger Causality</i>	176
8.3.4	<i>Nonlinear Granger Causality</i>	177
8.4	Copulas	179
8.4.1	<i>Fitting Copula Models</i>	180
8.4.2	<i>Parametric Copulas</i>	181
8.4.3	<i>Extending beyond Two Random Variables</i>	183
8.4.4	<i>Software</i>	185
8.5	Types of Dependence	185
8.5.1	<i>Positive and Negative Dependence</i>	185
8.5.2	<i>Tail Dependence</i>	187
	References	188
9	Correlated Poisson Processes and Their Applications in Financial Modeling	191
	<i>Alexander Kreinin</i>	
9.1	Introduction	191
9.2	Poisson Processes and Financial Scenarios	193
9.2.1	<i>Integrated Market–Credit Risk Modeling</i>	193
9.2.2	<i>Market Risk and Derivatives Pricing</i>	194
9.2.3	<i>Operational Risk Modeling</i>	194
9.2.4	<i>Correlation of Operational Events</i>	195
9.3	Common Shock Model and Randomization of Intensities	196
9.3.1	<i>Common Shock Model</i>	196
9.3.2	<i>Randomization of Intensities</i>	196
9.4	Simulation of Poisson Processes	197
9.4.1	<i>Forward Simulation</i>	197
9.4.2	<i>Backward Simulation</i>	200
9.5	Extreme Joint Distribution	207

9.5.1	<i>Reduction to Optimization Problem</i>	207
9.5.2	<i>Monotone Distributions</i>	208
9.5.3	<i>Computation of the Joint Distribution</i>	214
9.5.4	<i>On the Frechet–Hoeffding Theorem</i>	215
9.5.5	<i>Approximation of the Extreme Distributions</i>	217
9.6	<i>Numerical Results</i>	219
9.6.1	<i>Examples of the Support</i>	219
9.6.2	<i>Correlation Boundaries</i>	221
9.7	<i>Backward Simulation of the Poisson–Wiener Process</i>	222
9.8	<i>Concluding Remarks</i>	227
	<i>Acknowledgments</i>	228
Appendix A		229
A.1	<i>Proof of Lemmas 9.2 and 9.3</i>	229
A.1.1	<i>Proof of Lemma 9.2</i>	229
A.1.2	<i>Proof of Lemma 9.3</i>	230
	<i>References</i>	231
10	<i>CVaR Minimizations in Support Vector Machines</i>	233
	<i>Jun-ya Gotoh and Akiko Takeda</i>	
10.1	<i>What Is CVaR?</i>	234
10.1.1	<i>Definition and Interpretations</i>	234
10.1.2	<i>Basic Properties of CVaR</i>	238
10.1.3	<i>Minimization of CVaR</i>	240
10.2	<i>Support Vector Machines</i>	242
10.2.1	<i>Classification</i>	242
10.2.2	<i>Regression</i>	246
10.3	<i>ν-SVMs as CVaR Minimizations</i>	247
10.3.1	<i>ν-SVMs as CVaR Minimizations with Homogeneous Loss</i>	247
10.3.2	<i>ν-SVMs as CVaR Minimizations with Nonhomogeneous Loss</i>	251
10.3.3	<i>Refining the ν-Property</i>	253
10.4	<i>Duality</i>	256
10.4.1	<i>Binary Classification</i>	256
10.4.2	<i>Geometric Interpretation of ν-SVM</i>	257
10.4.3	<i>Geometric Interpretation of the Range of ν for ν-SVC</i>	258
10.4.4	<i>Regression</i>	259
10.4.5	<i>One-class Classification and SVDD</i>	259
10.5	<i>Extensions to Robust Optimization Modelings</i>	259
10.5.1	<i>Distributionally Robust Formulation</i>	259
10.5.2	<i>Measurement-wise Robust Formulation</i>	261
10.6	<i>Literature Review</i>	262
10.6.1	<i>CVaR as a Risk Measure</i>	263
10.6.2	<i>From CVaR Minimization to SVM</i>	263
10.6.3	<i>From SVM to CVaR Minimization</i>	263
10.6.4	<i>Beyond CVaR</i>	263
	<i>References</i>	264

11	Regression Models in Risk Management	266
	<i>Stan Uryasev</i>	
11.1	Introduction	267
11.2	Error and Deviation Measures	268
11.3	Risk Envelopes and Risk Identifiers	271
	<i>11.3.1 Examples of Deviation Measures \mathcal{D}, Corresponding Risk Envelopes \mathcal{Q}, and Sets of Risk Identifiers $\mathcal{QD}(X)$</i>	<i>272</i>
11.4	Error Decomposition in Regression	273
11.5	Least-Squares Linear Regression	275
11.6	Median Regression	277
11.7	Quantile Regression and Mixed Quantile Regression	281
11.8	Special Types of Linear Regression	283
11.9	Robust Regression	284
	References, Further Reading, and Bibliography	287
Index		289

1

Overview

Financial Signal Processing and Machine Learning

Ali N. Akansu¹, Sanjeev R. Kulkarni², and Dmitry Malioutov³

¹*New Jersey Institute of Technology, USA*

²*Princeton University, USA*

³*IBM T.J. Watson Research Center, USA*

1.1 Introduction

In the last decade, we have seen dramatic growth in applications for signal-processing and machine-learning techniques in many enterprise and industrial settings. Advertising, real estate, healthcare, e-commerce, and many other industries have been radically transformed by new processes and practices relying on collecting and analyzing data about operations, customers, competitors, new opportunities, and other aspects of business. The financial industry has been one of the early adopters, with a long history of applying sophisticated methods and models to analyze relevant data and make intelligent decisions – ranging from the quadratic programming formulation in Markowitz portfolio selection (Markowitz, 1952), factor analysis for equity modeling (Fama and French, 1993), stochastic differential equations for option pricing (Black and Scholes, 1973), stochastic volatility models in risk management (Engle, 1982; Hull and White, 1987), reinforcement learning for optimal trade execution (Bertsimas and Lo, 1998), and many other examples. While there is a great deal of overlap among techniques in machine learning, signal processing and financial econometrics, historically, there has been rather limited awareness and slow permeation of new ideas among these areas of research. For example, the ideas of stochastic volatility and copula modeling, which are quite central in financial econometrics, are less known in the signal-processing literature, and the concepts of sparse modeling and optimization that have had a transformative impact on signal processing and statistics have only started to propagate slowly into financial

applications. The aim of this book is to raise awareness of possible synergies and interactions among these disciplines, present some recent developments in signal processing and machine learning with applications in finance, and also facilitate interested experts in signal processing to learn more about applications and tools that have been developed and widely used by the financial community.

We start this chapter with a brief summary of basic concepts in finance and risk management that appear throughout the rest of the book. We present the underlying technical themes, including sparse learning, convex optimization, and non-Gaussian modeling, followed by brief overviews of the chapters in the book. Finally, we mention a number of highly relevant topics that have not been included in the volume due to lack of space.

1.2 A Bird's-Eye View of Finance

The financial ecosystem and markets have been transformed with the advent of new technologies where almost any financial product can be traded in the globally interconnected cyberspace of financial exchanges by anyone, anywhere, and anytime. This systemic change has placed real-time data acquisition and handling, low-latency communications technologies and services, and high-performance processing and automated decision making at the core of such complex systems. The industry has already coined the term *big data finance*, and it is interesting to see that technology is leading the financial industry as it has been in other sectors like e-commerce, internet multimedia, and wireless communications. In contrast, the knowledge base and exposure of the engineering community to the financial sector and its relevant activity have been quite limited. Recently, there have been an increasing number of publications by the engineering community in the finance literature, including *A Primer for Financial Engineering* (Akansu and Torun, 2015) and research contributions like Akansu *et al.*, (2012) and Pollak *et al.*, (2011). This volume facilitates that trend, and it is composed of chapter contributions on selected topics written by prominent researchers in quantitative finance and financial engineering.

We start by sketching a very broad-stroke view of the field of finance, its objectives, and its participants to put the chapters into context for readers with engineering expertise. Finance broadly deals with all aspects of money management, including borrowing and lending, transfer of money across continents, investment and price discovery, and asset and liability management by governments, corporations, and individuals. We focus specifically on trading where the main participants may be roughly classified into hedgers, investors, speculators, and market makers (and other intermediaries). Despite their different goals, all participants try to balance the two basic objectives in trading: to maximize future expected rewards (returns) and to minimize the risk of potential losses.

Naturally, one desires to buy a product cheap and sell it at a higher price in order to achieve the ultimate goal of profiting from this trading activity. Therefore, the expected return of an investment over any holding time (horizon) is one of the two fundamental performance metrics of a trade. The complementary metric is its variation, often measured as the standard deviation over a time window, and called investment risk or market risk.¹ Return and risk are two typically conflicting but interwoven measures, and risk-normalized return (Sharpe ratio)

¹ There are other types of risk, including credit risk, liquidity risk, model risk, and systemic risk, that may also need to be considered by market participants.

finds its common use in many areas of finance. Portfolio optimization involves balancing risk and reward to achieve investment objectives by optimally combining multiple financial instruments into a portfolio. The critical ingredient in forming portfolios is to characterize the statistical dependence between prices of various financial instruments in the portfolio. The celebrated Markowitz portfolio formulation (Markowitz, 1952) was the first principled mathematical framework to balance risk and reward based on the covariance matrix (also known as the variance-covariance or VCV matrix in finance) of returns (or log-returns) of financial instruments as a measure of statistical dependence. Portfolio management is a rich and active field, and many other formulations have been proposed, including risk parity portfolios (Roncalli, 2013), Black–Litterman portfolios (Black and Litterman, 1992), log-optimal portfolios (Cover and Ordentlich, 1996), and conditional value at risk (cVaR) and coherent risk measures for portfolios (Rockafellar and Uryasev, 2000) that address various aspects ranging from the difficulty of estimating the risk and return for large portfolios to the non-Gaussian nature of financial time series, and to more complex utility functions of investors.

The recognition of a price inefficiency is one of the crucial pieces of information to trade that product. If the price is deemed to be low based on some analysis (e.g. fundamental or statistical), an investor would like to buy it with the expectation that the price will go up in time. Similarly, one would shortsell it (borrow the product from a lender with some fee and sell it at the current market price) when its price is forecast to be higher than what it should be. Then, the investor would later buy to cover it (buy from the market and return the borrowed product back to the lender) when the price goes down. This set of transactions is the building block of any sophisticated financial trading activity. The main challenge is to identify price inefficiencies, also called *alpha* of a product, and swiftly act upon it for the purpose of making a profit from the trade. The efficient market hypothesis (EMH) stipulates that the market instantaneously aggregates and reflects all of the relevant information to price various securities; hence, it is impossible to beat the market. However, violations of the EMH assumptions abound: unequal availability of information, access to high-speed infrastructure, and various frictions and regulations in the market have fostered a vast and thriving trading industry.

Fundamental investors find *alpha* (i.e., predict the expected return) based on their knowledge of enterprise strategy, competitive advantage, aptitude of its leadership, economic and political developments, and future outlook. Traders often find inefficiencies that arise due to the complexity of market operations. Inefficiencies come from various sources such as market regulations, complexity of exchange operations, varying latency, private sources of information, and complex statistical considerations. An *arbitrage* is a typically short-lived market anomaly where the same financial instrument can be bought at one venue (exchange) for a lower price than it can be simultaneously sold at another venue. Relative value strategies recognize that similar instruments can exhibit significant (unjustified) price differences. Statistical trading strategies, including statistical arbitrage, find patterns and correlations in historical trading data using machine-learning methods and tools like factor models, and attempt to exploit them hoping that these relations will persist in the future. Some market inefficiencies arise due to unequal access to information, or the speed of dissemination of this information. The various sources of market inefficiencies give rise to trading strategies at different frequencies, from high-frequency traders who hold their positions on the order of milliseconds, to midfrequency trading that ranges from intraday (holding no overnight position) to a span of a few days, and to long-term trading ranging from a few weeks to years. High-frequency trading requires state-of-the-art computing, network communications, and