


The Elements of Quantitative Investing PDF

Visit the link below to download the full version of the ebook

[DOWNLOAD NOW](#)

GIUSEPPE A. PALEOLOGO

THE ELEMENTS OF QUANTITATIVE INVESTING




Scan to Download
or Type the Link

ebook.ac/elements

GIUSEPPE A. PALEOLOGO

THE ELEMENTS OF QUANTITATIVE INVESTING



WILEY

The Elements of Quantitative Investing

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation and financial instrument analysis, as well as much more.

For a list of available titles, visit our Web site at www.WileyFinance.com.

The Elements of Quantitative Investing

Giuseppe A. Paleologo

WILEY

Copyright © 2025 by Giuseppe A. Paleologo. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

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, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: Wiley and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. 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. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data Applied for:

Print ISBN: 9781394265459

ePDF ISBN: 9781394265473

epub ISBN: 9781394265466

Cover Image: Courtesy of Giuseppe A. Paleologo

Cover Design: Wiley

To Tofu, again

Contents

Acknowledgments	xv
Introduction	xvii
Notation	xxiii
Chapter 1	
The Map and the Territory	1
1.1 The Securities	3
1.2 Modes of Exchange	5
1.3 Who Are the Market Participants?	6
1.3.1 The Sell Side	6
1.3.2 The Buy Side	9
1.4 Where Do Excess Returns Come From?	12
1.5 The Elements of Quantitative Investing	15
Chapter 2	
Univariate Returns	20
2.1 Returns	21
2.1.1 Definitions	21
2.1.2 Excess Returns	23
2.1.3 Log Returns	23

	2.1.4	Estimating Prices and Returns	24
	2.1.5	Stylized Facts	26
2.2		Conditional Heteroskedastic Models	30
	2.2.1	GARCH(1, 1) and Return Stylized Facts	32
	2.2.2	GARCH as Random Recursive Equations	34
	2.2.3	★GARCH(1, 1) Estimation	36
	2.2.4	Realized Volatility	37
2.3		State-Space Estimation of Variance	40
	2.3.1	Muth's Original Model: EWMA	40
	2.3.2	★The Harvey–Shephard Model	44
2.4		★Appendix	46
	2.4.1	The Kalman Filter	46
	2.4.2	Kalman Filter Examples	49
2.5		Exercises	51
Chapter 3		Interlude: What Is Performance?	53
	3.1	Expected Return	54
	3.2	Volatility	54
	3.3	Sharpe Ratio	55
	3.4	Capacity	58
Chapter 4		Linear Models of Returns	61
	4.1	Factor Models	62
	4.2	Interpretations of Factor Models	65
		4.2.1 Graphical Model	66
		4.2.2 Superposition of Effects	66
		4.2.3 Single-Asset Product	67
	4.3	Alpha Spanned and Alpha Orthogonal	68
	4.4	Transformations	71
		4.4.1 Rotations	71
		4.4.2 Projections	73
		4.4.3 Push-Outs	74
	4.5	Applications	75
		4.5.1 Performance Attribution	75
		4.5.2 Risk Management: Forecast and Decomposition	76

	4.5.3	Portfolio Management	80
	4.5.4	Alpha Research	80
	4.6	Factor Models Types	81
	4.7	★Appendix	82
	4.7.1	Linear Regression	82
	4.7.2	Linear Regression Decomposition	86
	4.7.3	The Frisch–Waugh–Lovell Theorem	87
	4.7.4	The Singular Value Decomposition	89
	4.8	Exercises	92
Chapter 5		Evaluating Risk	94
	5.1	Evaluating the Covariance Matrix	95
	5.1.1	Robust Loss Functions for Volatility Estimation	95
	5.1.2	Application to Multivariate Returns	97
	5.2	Evaluating the Precision Matrix	100
	5.2.1	Minimum-Variance Portfolios	100
	5.2.2	Mahalanobis Distance	101
	5.3	Ancillary Tests	102
	5.3.1	Model Turnover	103
	5.3.2	Testing Betas	103
	5.3.3	Coefficient of Determination?	104
	5.4	★Appendix	107
	5.4.1	Proof for Minimum-Variance Portfolios	107
Chapter 6		Fundamental Factor Models	110
	6.1	The Inputs and the Process	111
	6.1.1	The Inputs	111
	6.1.2	The Process	114
	6.2	Cross-Sectional Regression	115
	6.2.1	Rank-Deficient Loadings Matrices	118
	6.3	Estimating the Factor Covariance Matrix	120
	6.3.1	Factor Covariance Matrix Shrinkage	121
	6.3.2	Dynamic Conditional Correlation	122

	6.3.3	Short-Term Volatility Updating	122
	6.3.4	Correcting for Autocorrelation in Factor Returns	124
6.4		Estimating the Idiosyncratic Covariance Matrix	125
	6.4.1	Exponential Weighting	125
	6.4.2	Visual Inspection	125
	6.4.3	Short-Term Idio Update	126
	6.4.4	Off-Diagonal Clustering	127
	6.4.5	Idiosyncratic Covariance Matrix Shrinkage	131
6.5		Winsorization of Returns	131
6.6		★Advanced Model Topics	133
	6.6.1	Linking Models	133
	6.6.2	Currency Rebasing	139
6.7		A Tour of Factors	141
Chapter 7		Statistical Factor Models	147
7.1		Statistical Models: The Basics	149
	7.1.1	Best Low-Rank Approximation and PCA	149
	7.1.2	Maximum Likelihood Estimation and PCA	152
	7.1.3	Cross-Sectional and Time-Series Regressions via SVD	155
7.2		Beyond the Basics	155
	7.2.1	The Spiked Covariance Model	156
	7.2.2	Spectral Limit Behavior of the Spiked Covariance Model	158
	7.2.3	Optimal Shrinkage of Eigenvalues	160
	7.2.4	Eigenvalues: Experiments versus Theory	162
	7.2.5	Choosing the Number of Factors	162
7.3		Real-Life Stylized Behavior of PCA	165
	7.3.1	Concentration of Eigenvalues	166
	7.3.2	Controlling the Turnover of Eigenvectors	168

7.4	Interpreting Principal Components	173
7.4.1	The Clustering View	173
7.4.2	The Regression View	174
7.5	Statistical Model Estimation in Practice	176
7.5.1	Weighted and Two-Stage PCA	176
7.5.2	Implementing Statistical Models in Production	179
7.6	★Appendix	181
7.6.1	Exercises and Extensions to PCA	181
7.6.2	Asymptotic Properties of PCA	185
Chapter 8	Evaluating Excess Returns	188
8.1	Backtesting Best Practices	190
8.1.1	Data Sourcing	190
8.1.2	Research Process	191
8.2	The Backtesting Protocol	195
8.2.1	Cross-Validation and Walk-Forward	195
8.3	The Rademacher Anti-Serum (RAS)	200
8.3.1	Setup	200
8.3.2	Main Result and Interpretation	203
8.4	Some Empirical Results	208
8.4.1	Simulations	208
8.4.2	Historical Anomalies	210
8.5	★Appendix	214
8.5.1	Proofs for RAS	214
Chapter 9	Portfolio Management: The Basics	220
9.1	Why Mean-Variance Optimization?	221
9.2	Mean-Variance Optimal Portfolios	223
9.3	Trading in Factor Space	229
9.3.1	Factor-Mimicking Portfolios	229
9.3.2	Adding, Estimating, and Trading a New Factor	232
9.3.3	Factor Portfolios from Sorts?	235
9.4	Trading in Idio Space	236
9.5	Drivers of Information Ratio: Information Coefficient and Diversification	237
9.6	Aggregation: Signals versus Portfolios	240

	9.7	★Appendix	244
	9.7.1	Some Useful Results from Linear Algebra	244
	9.7.2	Some Portfolio Optimization Problems	245
	9.7.3	Optimality of FMPs	245
	9.7.4	Single-Factor Covariance Matrix Updating	247
Chapter 10		Beyond Simple Mean-Variance	250
	10.1	Shortcomings of Naïve MVO	251
	10.2	Constraints and Modified Objectives	254
	10.2.1	Types of Constraints	257
	10.2.2	Do Constraints Improve or Worsen Performance?	261
	10.2.3	Constraints as Penalties	261
	10.3	How Does Estimation Error Affect the Sharpe Ratio?	267
	10.3.1	The Impact of Alpha Error	268
	10.3.2	The Impact of Risk Error	269
	10.4	★Appendix	270
	10.4.1	Theorems on Sharpe Efficiency Loss	270
Chapter 11		Market-Impact-Aware Portfolio Management	276
	11.1	Market Impact	277
	11.1.1	Temporary Market Impact	278
	11.2	Finite-Horizon Optimization	284
	11.3	Infinite-Horizon Optimization	286
	11.3.1	Comparison to Single-Period Optimization	289
	11.3.2	The No-Market-Impact Limit	290
	11.3.3	Optimal Liquidation	290
	11.3.4	Deterministic Alpha	291
	11.3.5	AR(1) Signal	291
	11.4	★Appendix	293
	11.4.1	Proof of the Infinite-Horizon Quadratic Problem	293

Chapter 12	Hedging	297
	12.1 Toy Story	298
	12.2 Factor Hedging	301
	12.2.1 The General Case	301
	12.3 Hedging Tradeable Factors with Time-Series Betas	304
	12.4 Factor-Mimicking Portfolios of Time Series	307
	12.5 ★Appendix	309
Chapter 13	Dynamic Risk Allocation	312
	13.1 The Kelly Criterion	314
	13.2 Mathematical Properties	321
	13.3 The Fractional Kelly Strategy	323
	13.4 Fractional Kelly and Drawdown Control	327
Chapter 14	Ex-Post Performance Attribution	333
	14.1 Performance Attribution: The Basics	335
	14.2 Performance Attribution with Errors	336
	14.2.1 Two Paradoxes	336
	14.2.2 Estimating Attribution Errors	337
	14.2.3 Paradox Resolution	339
	14.3 Maximal Performance Attribution	340
	14.4 Selection versus Sizing Attribution	347
	14.4.1 Connection to the Fundamental Law of Active Management	351
	14.4.2 Long–Short Performance Attribution	351
	14.5 Appendix★	352
	14.5.1 Proof of the Selection versus Sizing Decomposition	352
Chapter 15	A Coda about Leitmotifs	357
	References	359
	Index	373

Acknowledgments

The following people have proofread, commented on, discussed or supported the development of this book. Some of them are old friends, most are strangers who spontaneously offered to help. Meeting them (usually on X or on LinkedIn) has been an unforeseen pleasure in the writing of the book, and I am grateful for their generosity and the many corrections they sent my way. Unfortunately all the remaining errors are my responsibility.³

@0xdf, Rashad Ahmed, Amir Aliev, Sandro C. Andrade, Samuel Babichenko, Ashish Bajpaj, Soumyadipta Banerjee, Rayan Ben Redjeb, Jerome Benveniste, Igor Berman, Jon Beyer, Victor Bomers, David Bonnerot, Thomas Byrne, Ben Chaddha, Yuyao Che, Leif Cussein, Vladimir Cvetkovic, Drake Daly, Alex Darby, Ruolag Deng, Jane Doe, John Doe (Jane's occasionally hostile brother), Saarialho Eero, Bill Falloon, Frank Fan, Marta Filizola, Tom Fleming, Claudio Fontana, Fan Gao, Ernesto Guridi, Touko Haapanen, Victor Haghani, Leila Hardy, Levon Haykazyan, Erik Hellgren, Alex Hermeneanu, Leonhard Hochfilzer, Clint Howard, Jack Huang, Kevin Jacobs,

³ And my cat's.

Fredrik Jäfvert, Bob Jansen, Amanda Jiménez, Yassine Kachrad, Abhinav R. Kannan, Shahin Khobahi, D. K., Marco Kuhlen, Markku Kurtti, Kofi Kwapong, Jessie Li, Vitoria Lima, @jesse_livermore, Samuel Londner, Dustin Lorshbough, Marco Lucchesi, Joseph Maestri, Baridhi Malakar, Emanuel Malek, Paul-Henri Memin, Francesco Mina, Simon Minovitsky, Caroline Moon, Roberto Moura, Abhijit Naik, Graeme Newby, Juntunen Nikolaus, Mikhail Novichkov, Andrew Novotny, Adam Nunes, Tom Ó Nualláin, Yogesh Padmanabhan, Jeff Park, Damon Petersen, Henrik Hiro Pettersson, Saharat Phomthong, Alexis Plascencia, Gabriele Pompa, David Popovic, Akshay Prasad, Miika Purola, Siddhant Pyasi, Cas Rijnierse, Viktor Salihu, Mateus Sampaio, Ali Sanjari, Douglas Ricardo Sansão, Tanmay Satpathy, Guilherme Saturnino, Abishek Saxena, Christian Schmitz, Maarten Scholl, Timothy Shchetilin, Shawn Sheng, Nakul Shenoy, Marco Signoretto, Ashutosh Singh, Matt Smith, Richárd Süveges, @systematicls, Samer Takriti, Gianmario Tamagnone, Xiaolong Tan, Kaspar Thommen, Maxim Tishin, Shahpour Turkian, Nathan Ueda, Carlos Ungil, Sinan Unver, Ricardo Valdez, Derrick VanGennep, Paul Vriend, Jay Vyas, Andrew Wang, Shida Wang, Simon Xiang, Jeremy Young, Alvin Zhang, Howard Zhang, Siqi Zheng.

I owe special thanks to Brandon DiNunno, Denis Dmitriev, Jean-François Fortin, Antoine Liutkus, Dustin Lorshbough, Tom Mainiero, Krutarth Satoskar, and Trent Spears for their deep and wide reading of the manuscript.

G.A.P.

Introduction

This book originates from notes I wrote for two university courses. The first is ORIE5256 – Topics in Risk Management and Portfolio Construction, a course offered in the program for M.S. in Financial Engineering at Cornell University. The second is MATH-GA 2708.001 – Algorithmic Trading & Quantitative Strategies, offered in the Mathematics Department at New York University. When I set out to write this book, my objective was to write the quantitative introduction I had wanted to read at the beginning of my journey in finance. Given the scope and goals of quantitative investing, it is only possible to cover a small fraction of it in a course, or even in a book. To address this problem, I made three choices.

First and most important, I aim for synthesis. A book is, first of all, a knowledge filter. In the preface to his classic (Kelley, 1955), Kelley wrote that he wanted to title his book “What Every Young Analyst Should Know”; that book was barely three hundred pages long. It still feels fresh and necessary today. In order to keep my book of manageable length, my working principle has been to focus on real-world problems and then use the *simplest* techniques that allow me to address the problem at hand. A recurrent theme in the book is that almost everything in it is either linear

or quadratic. In the process of writing, I have ruthlessly eliminated topics of secondary importance, material that was too hard for the payoff that it gave the readers, and also topics or ideas that are not sufficiently well-formed, or too experimental. Even if you choose not to read my book, I implore you to internalize the following lesson, learned by practitioners through sweat and tears: *theory is cheap*. There are thousands and thousands of theory papers, in love with technical virtuosity but oblivious of reality. Do not fall into temptation; by applications be driven.¹

Second, I consider risk management and portfolio management as intrinsically connected. Asset return modeling, volatility estimation, portfolio optimization, *ex-ante* and *ex-post* performance analytics are all linked. For example, hedging belongs to risk and portfolio optimization, and analysis of performance feeds back into portfolio construction. I have avoided redundancy as much as possible. Sections often refer to earlier ones or are linked to later ones. As I was revising my book draft, whenever I found I had introduced some topic (often because I was infatuated with it) and then never used it, I exiled it to a long file made out of “rejected sections” that lives in my laptop. That is the sad part. The happy part: it’s surprising how tall a tree can grow with a bit of pruning. Out of metaphor, there is a lot of material in this book, and it gets challenging at times.

Third, I occasionally integrate some standard financial results approaches with tools from the field of statistical learning. The former is applied in fundamental factor modeling, portfolio optimization, and performance attribution. I use the latter for the estimation of statistical models and backtesting. My hope is that the integration of these different approaches is seamless.

The questions that I address in this book are:

1. How do I model returns in a way that allows me to generate risk and return forecasts?
 - (a) What are excess returns?
 - (b) How do I model multivariate returns?
 - (c) How do I describe and forecast risk?
 - (d) How do I test risk forecasts and return forecasts?
 - (e) How do I backtest alphas?

¹ References to “By Demons be Driven” by Pantera and to Warcraft 3 are intentional.

2. How do I monetize these signals?
 - (a) How do I optimize a portfolio?
 - (b) What is the impact of risk and alpha errors on performance?
 - (c) How do I account for transaction costs in portfolio management?
 - (d) How do I hedge a portfolio?
3. How do I improve?
 - (a) How do I allocate risk over time?
 - (b) How do I distinguish skill from luck?

The style of the book is also, I hope, a bit different. I have kept in mind the six values that Italo Calvino (Calvino, 1999) hoped to preserve in the current millennium: Lightness, Quickness, Exactitude, Visibility, Multiplicity, and Consistency. My aversion to advanced mathematics notwithstanding, I must warn the reader that the book is not easy. During my lectures, I have induced more than one student into a comatose stupor. Afterwards, a few students left finance altogether and successfully pursued careers in entertainment. Another student keeps sending me postcards from Ibiza. A handful have become portfolio managers at hedge funds and risk managers. Yet, it is the easiest book I could write for the task at hand, and it is written in the *friendliest* style I am capable of. Also, I would be lying—and conveying the wrong message—if I claimed that the problems I present are now settled, and that the book is the last word on the subject. On the contrary, you and I are in this book together, and together we shall keep a beginner’s mind (Suzuki, 1970): a spirit of openness and curiosity, even when facing advanced topics. I will point out the limits of my theories and the open problems for you to work on. If you are old enough to have lived in the seventies and liked punk music, you may remember a cyclostyled zine (Figure 1). On its second page it showed three open chords; below them, a command: “NOW FORM A BAND.” May this book be your field guide to being a punk quantitative researcher. It will be a life well lived.

Prerequisites

The book should be accessible to a beginning graduate or advanced undergraduate student in Physics, Mathematics, Statistics, or Engineering. This means having a working relationship, and if possible a romantic one, with advanced linear algebra, probability theory, and statistics. Even more important is to have a deep interest in quantitative modeling of real-life




Figure 1 Punk fanzine *Sideburn* #1, page 2 (1977). Source: Flickr.com/Dunk.

phenomena. Many readers will be either members of a systematic trading team, or work as quantitative researchers in the central team of a hedge fund or a quantitative asset manager.

The book's material is organized in such a way that you do not need to go through mathematical proofs. You can rely only on informal statements of mathematical results in the main body of the chapters, and that will suffice to understand the main points. The appendices at the end of the chapters contain more rigorous statements, proofs, and background material. If you plan on actively doing research, you should study them, eventually.

Even if you read only the main body, you should be used to thinking in mathematical models. The *Book of Nature* is written in a mathematical language.² Be comfortable with:

² "Philosophy is written in this grand book, the universe, which stands continually open to our gaze. But the book cannot be understood unless one first learns to comprehend the language and read the letters in which it is composed. It is written in the language of mathematics, and its characters are triangles, circles, and other geometric figures without which it is humanly impossible to understand a single word of it; without these, one wanders about in a dark labyrinth" (Galilei, 1623).

- Working with linear algebra, at least at the level of Strang (2019) and Trefethen and Bau (1997).
- Some applied probability, at the level of Ross (2023). Exposure to some simple control theory and state-space models helps. You can come to this from econometrics (Harvey, 1990; Shumway and Stoffer, 2011), control theory (Simon2006), or statistics (Hyndman et al., 2008).
- Some optimization modeling is a plus. The first few chapters of Boyd and Vandenberghe (2004) would be ideal. However, I will cover the basic theory in an appendix.

Organization

Like Caesar's Gaul, the book is broadly divided into three parts. The first part focuses on returns modeling. I cover the basics of GARCH early on because they are needed for factor modeling, and then I cover factor models because they are necessary for everything. I have separate chapters for fundamental and statistical models. These topics are covered in depth, and both the treatment and some of the modeling approaches are novel. Finally, I cover data snooping/backtesting as a separate chapter, since it is a central element of the investment process.

The second part is devoted to portfolio construction and performance analysis, both *ex ante* and *ex post*. The focus is on mean-variance optimization (MVO). I emphasize the geometric intuition behind much of mean-variance optimization. Rotations, projections, and angles are prevalent throughout. This allows for a synthetic, elegant characterization of performance and for concise proofs. The statistics of the Sharpe Ratio are covered in some detail. The decomposition of payoffs into timing components, factor and idiosyncratic Profit and Loss (PnL), and stock selection versus sizing of positions is rigorously demonstrated. Model error plays an important role in this part. If an optimization problem is Othello, then model error must be Iago: it can drive the optimization insane. Unlike in Shakespeare's tragedies, we can try to rewrite the endings and turn them into comedies.

The third part is the shortest. It contains results about intertemporal volatility allocation and performance attribution. These are essential components of the investment process and belong in a book with the word "Elements" in its title.




Figure 2 Chapter dependencies.

Each chapter is organized like an onion. The first sections convey the essential ideas using simple quantitative methods and are aimed at a broad audience. Sections marked with a star “★” are more advanced and can be skipped on a first reading. Proofs of new results or basic technical material are relegated to the appendices at the end of the chapters. The goal is not to disrupt the flow of learning. As mentioned at the beginning of this preface, the content of this book was taught first and written later. It is meant to be read aloud and discussed, but it should be suitable for self-study. The dependencies among the chapters are shown in Figure 2.

Giuseppe “gappy” A. Paleologo
 Riverdale, New York
 March 21, 2025

Notation

$\mathbb{R}$	field of real numbers
$\mathbb{R}^+$	$[0, \infty)$
$\mathbb{N}$	set of natural numbers
$x, y, \dots$	scalars
$\mathbf{x}, \mathbf{y}, \dots$	vectors (assumed to be $n \times 1$)
$\mathbf{X}, \mathbf{Y}, \dots$	matrices
$\mathbf{x}^\top, \mathbf{X}^\top$	vector or matrix transpose
$\mathbf{X}^+$	Moore–Penrose pseudo-inverse
$\mathbf{\Omega}$	covariance matrix
$\text{diag}(x_1, \dots, x_n)$	diagonal matrix with scalars $x_1, \dots, x_n$ on the main diagonal
$[\mathbf{x}]_i, x_i$	i th element of a vector
$[\mathbf{X}]_{i,j}$	element of a matrix on i th row and j th column
$[\mathbf{X}]_{i\cdot}$ or $\mathbf{x}_i$	i th row of matrix X
$[\mathbf{X}]_{\cdot,j}$ or $\mathbf{x}^j$	j th column of matrix X
trace $(\mathbf{X})$	trace operator: $\text{trace}(\mathbf{X}) = \sum_i [\mathbf{X}]_{i,i}$

$ \mathbf{X} $	determinant of a matrix
$\delta_{i,j}$	Kroneker's delta: $\delta_{i,j} = 1$ if $i = j$, 0 otherwise
$\delta(t)$	Dirac's delta function: $\delta(t) = 0$ for $t \neq 0$, and $\int_{-\infty}^{\infty} \delta(t) dt = 1$
$\mathbf{e}$	vector whose elements are all ones: $[\mathbf{e}]_i = 1$ for all i
$\mathbf{1}\{x\}$	indicator function, equal to 1 if x is true, 0 otherwise
$\mathbf{e}_i$	vector whose i th element is 1 and the others are zero: $[\mathbf{e}_i]_j = \delta_{i,j}$
$\mathbf{I}_m$	identity matrix of size $m \times m$
$\ \mathbf{x}\ _p$	p -norm, for $p \geq 1$: $(\sum_{i=1}^n x_i ^p)^{1/p}$
$\ \mathbf{x}\ $	2-norm of a vector
$\ \mathbf{x}\ _{\mathbf{Q}}$	$\sqrt{\mathbf{x}^T \mathbf{Q} \mathbf{x}}$, for a positive-definite matrix $\mathbf{Q}$
$\ \mathbf{H}\ $	operator norm for a positive-definite matrix $\mathbf{H}$: $\max_{\mathbf{x}} \ \mathbf{H}\mathbf{x}\ / \ \mathbf{x}\ := \left(\max_{\mathbf{x}} \mathbf{x}^T (\mathbf{H}^T \mathbf{H}) \mathbf{x} / \ \mathbf{x}\ ^2 \right)^{1/2}$
$\lfloor x \rfloor$	largest integer that is less than, or equal to, x
$\langle \mathbf{x}, \mathbf{y} \rangle$	scalar product of two vectors, i.e., $\mathbf{x}^T \mathbf{y}$
$\langle \mathbf{A}, \mathbf{B} \rangle$	scalar product of two matrices: $\langle \mathbf{A}, \mathbf{B} \rangle \in \mathbb{R} := \text{trace}(\mathbf{A}^T \mathbf{B})$
$S_C(\mathbf{a}, \mathbf{b})$	cosine similarity $S_C(\mathbf{x}, \mathbf{y}) := \langle \mathbf{x}, \mathbf{y} \rangle / (\ \mathbf{a}\ \ \mathbf{b}\)$
$\mathbf{x} \circ \mathbf{y} \in \mathbb{R}^n$,	Hadamard (element-wise) product of two vectors (matrices):
$\mathbf{X} \circ \mathbf{Y} \in \mathbb{R}^{m \times n}$	$[\mathbf{x} \circ \mathbf{y}]_i := x_i y_i$, $[\mathbf{X} \circ \mathbf{Y}]_{i,j} := [\mathbf{X}]_{i,j} [\mathbf{Y}]_{i,j}$
$\xi \perp \zeta$	random variables ξ and ζ are independent
$E(\xi), E(\xi)$	expectation of random variables or random vectors
$E_{\xi}(f(\xi))$	expectation of a function f of random variable ξ
$\hat{E}(\mathbf{x})$	average of a vector $\mathbf{x}$: $\hat{E}(\mathbf{x}) := n^{-1} \sum_i x_i$
$\text{var}(x)$	variance of a random variable x
$\mathbf{L}_1 : \mathbb{R}^n \rightarrow \mathbb{R}^n$	shift operator: $L_1 \mathbf{x} := (x_2, x_3, \dots, x_n, 0)^T$, i.e., $[\mathbf{L}_1]_{i,j} := \delta_{i,j-1}$
rv	random variables
iid rv	independent, identically distributed random variables
$\stackrel{d}{=}$	equality in distribution of two random variables
$\sim$	rv distributed according to a distribution, e.g., $\xi \sim N(0, 1)$
$\asymp$	of the same order: $n \asymp m$ if $an < m < bn$ for some $0 < a \leq b$
$\Rightarrow$	convergence in distribution
$\equiv$	equivalence of optimization problems: for constants $a > 0$, $b \in \mathbb{R}$ $\max_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) \equiv \max_{\mathbf{x} \in \mathcal{S}} [af(\mathbf{x}) + b]$

Chapter 1

The Map and the Territory

The Questions

1. What are the essential components of quantitative investing?
2. What types of securities are involved and how are they traded?
3. Who are the main market participants and what roles do they play?
4. Where do excess returns in investing come from?
5. What are the key elements that form the analytical framework of a quantitative portfolio manager?

This chapter is a guide to the essential components of quantitative investing. When considering the meaning of a word, it's often instructive to go back to its etymology, so let's play this game. Despite being a

Germanic language, English adopted many words from Latin, sometimes by way of French. “Investing” comes from “Investire”, which in Latin meant “to cover with a vest”, or “to put in a vest”. So it should be hardly surprising that two thousand years later, vests would become the favorite garment of hedge fund managers. In the Middle Ages, the verb took on the additional meaning “to surround, to have ownership of”. It is also possible that the modern meaning overtook the old because, in ceremonies in which ownership was transferred, the new owner was “invested” with a cloak and other regalia. In Italian—the direct successor to Latin—the old meaning is gone, and “investire” only means “to receive possession of something”. As for “quantitative”, that is Latin too: “quantum”, a noun denoting something that can be measured, increased, and decreased. We will deal with ownership, sold and bought in units that can be measured, increased and decreased. This is, unfortunately, the whole of finance. You can own a house, a painting, a bet on the survival of humankind, or even an idea. Each one of these investment topics deserves its own book, written by a competent author. In writing this book, I have chosen to trade off generality in favor of detail. I have covered each subject with the goal in mind that you would have sufficient information to understand it, implement it, and critique it. However, even an analytical book needs an introduction that puts things in their proper context. In this chapter, I aim to provide that context. You will have a broad understanding of the classes of securities to which these methods apply; and of the way these securities are traded, and by whom. This is a necessary prerequisite to explore fundamental questions: Where are excess returns coming from? What causes these trading opportunities? Finally, I will present the essential components that make up the analytical framework of a quantitative portfolio manager. The underlying message is that to be successful, an investor must understand how things work. A seminal early book on investing is titled “The Intelligent Investor” (Graham, 2006). To double down on Latin, the original meaning of “intelligent” is “to read into something”, similar to “insightful” in the English language. Your success will come from reasoning about the behavior of your counterparties, the rules governing the trading of your assets, and the functioning of exchanges. Many budding quants focus on quantitative methods. The fact is that theory is cheap and is often not hard. What is hard is putting the right tool at the service of the right insight.

Finally, this is the only chapter without mathematics. You should enjoy it while it lasts.

1.1 The Securities

We will be concerned with *standardized products* that are *liquid*. We explain these concepts in more detail.

To “own” an object is effectively to own *claims* on that object in the future. If you own a house, you can live in it or rent it out (your claim) and it is yours. This claim is not absolute, however. In most countries, the local or central government may need your property for reasons of public welfare and can require you to exchange your claim for cash at a fair price. If you own a painting, you may enjoy it in the confines of your house, but may not necessarily own its reproduction rights. If you own a bet on the future of humanity, your counterparty may have some *force majeure* clauses that prevent it from paying (e.g., consider a zombie apocalypse scenario). Defining ownership of an “idea” is especially challenging and prone to be treated on an *ad hoc* basis. Compared to the infinite and ever-changing nature of the meaning of property rights, our coverage is very narrow. Specifically, we focus on the subset of contracts that are standardized and liquid. We buy and sell *standardized* claims. These claims come in a few varieties, and their attributes are clearly defined and known to all potential buyers and sellers. Examples are:

- *Equities and Exchange-Traded Funds (ETFs)*. These give us partial ownership in companies, or groups of companies, and entitle us to receive future cash payments generated by the economic activities of these companies.
- *Futures*. These contracts deliver a physical commodity or a cash payment contingent on the state of the world at a future date, at a price determined today.
- *Bonds*. These are contracts that allow the transfer of debt claims among parties. An investor lends money to a borrower, in exchange for a fixed cash flow in the future (e.g., periodic interest payments and a final payment). A bond makes this claim transferable to other lenders.
- *Vanilla options*. These are claims that depend on the future value of some underlying asset; for example, you may receive the right (but not the

obligation) to buy a stock at a future date, at a price determined today. The nature of these claims is standardized, hence the term “vanilla”.

- *Interest Rate Swaps (IRSs)*. These contracts allow the exchange of a certain, deterministic cash flow stream for an uncertain one, which depends on interest rates at future dates.
- *Credit Default Swaps (CDSs)*. These contracts insure the buyer against the failure of a company at a future date, in exchange for recurring fixed payments.

Further these contracts are *liquid*. For our purposes, a liquid contract is one that can be bought and sold at large enough sizes, and at sufficiently short time horizons, to enable quantitative strategies to be implemented. This means that if we plan to buy or sell a contract, we should be able to do so without incurring a transaction cost so high that our strategy is not economically attractive even for small trading sizes, and that the waiting time due to searching for a counterparty should not be so long as to make the transaction economically unattractive.

The properties of standardization and liquidity are closely intertwined. Increased standardization tends to enhance liquidity by consolidating demand, as it aggregates dispersed demand from bespoke products toward a smaller set of standardized ones. Furthermore, standardization streamlines the trading process, reducing transaction costs and simplifying contracts, thereby fostering investor trust and attracting more participants, thus enhancing liquidity. However, the downside is that customers may sacrifice the ability to trade certain useful product characteristics. Determining the optimal level of customization, even at the expense of liquidity, remains an ongoing process of learning and adaptation. For instance, prior to the 2008 financial crisis, CDSs exhibited greater variety. However, the “Big Bang” initiated by the International Swaps and Derivatives Association (ISDA) on April 8, 2009, simplified contract terms, including standardizing coupon rates (100 bps and 500 bps) and introducing a standard upfront payment, which played a pivotal role in restoring confidence in this asset class (Vause, 2010).

Trading and liquidity are at the core of the book. In order to better understand the trading process and the nature of liquidity, we should describe in some detail how trading on-exchange and over-the-counter happens.

1.2 Modes of Exchange

At any given time, economic agents want to buy or sell contracts. They want to do so quickly, securely, and cheaply. The three options around which trading is currently organized are exchanges, over-the-counter, and dark pools. Exchanges are venues in which the orders of buyers and sellers are anonymized and matched against each other. Orders are characterized by size, the number of contracts, direction (buy or sell), and price. They represent requests to buy or sell a number of contracts. The exchange records such active orders on a ledger, known as the limit-order book (LOB), and employs a set of priority rules to match buy and sell orders in the exchanges—aptly named a matching engine. In order to trade on an exchange, one must be a member of that exchange. Membership entails apparent benefits and less-apparent responsibilities. Market participants must maintain sound governance, risk processes, and capital structure.

Exchanges evolve continuously due to two driving forces. On one side, there is a push toward consolidation, which reduces operating costs and gives the owner pricing power. On the other side, technical and process innovations introduce new competitors into the market. In the United States alone, there are more than a dozen equity stock exchanges. Exchange-traded assets, such as stocks, options, and future contracts (including Forex Futures), are often liquid, although this condition is neither necessary nor sufficient: some exchange-traded assets are traded in minimal volumes and, therefore, are not liquid, and some very liquid products are traded off-exchange.

Other assets are not traded on exchanges, but *over-the-counter (OTC)*. In this case, the buyer or seller transacts through an institutional market participant, the broker-dealer, which is connected to other broker-dealers and facilitates the matching of orders. Bonds, IRSs, Forex spot currencies, Forex Forwards, and CDSs are examples of contracts traded OTC. Some of these, like currencies, are among the world's most liquid contracts. A precondition for liquidity is standardization. Think of a house. "The New York housing market" is very different from the stock market in that each of the 15.4 billion outstanding Apple shares (as of July 2024) is indistinguishable from the other and sells in a matter of seconds. In contrast, a house has many attributes that make it unique: location, size, age, blueprint, and condition. Another characteristic of liquid markets is the large number of participants. When numerous participants are involved in

a market, competing for a relatively low number of contracts, transactions become more frequent, and the necessity for bilateral bargaining diminishes. The ability of any individual participant to influence the price is significantly reduced. To illustrate, consider the housing market as a counterpoint: when selling a house, you typically negotiate with one specific buyer (out of a few eligible ones), who may spend many hours searching for the right property and may engage in intense bargaining, sometimes to the point of contention, to secure the best possible price.

Finally, *Dark Pools* (a type of ATS, or Alternative Trading System, which does not make its LOB transparent) are additional venues that are distinct from exchanges (although sometimes owned by them). Dark Pools address the needs of certain institutional investors to execute orders without displaying their trading intentions. By design, Dark Pools hide order details and only make trades details available after execution. As of 2024, approximately 16% of U.S. shares are traded on Dark Pools.

1.3 Who Are the Market Participants?

It is convention to partition traders into the *sell side* and the *buy side*. The former facilitates trading by providing services; the latter receives trade for their own benefit. Below I describe the participant types. For a more detailed description, see Harris (2003).

1.3.1 The Sell Side

The sell side comprises brokers, dealers, and broker-dealers.

- *Dealers*¹ fulfill their clients' demands, thus providing liquidity. They take the opposite side of the trade; they are profitable if, on average, they sell (buy) at a price higher (lower) than what they initially paid to buy (sell) the asset. The difference between buy and sell prices is the *spread*. When dealers interact with clients, they quote the buy price (the *bid price*) and the sell price (the *ask price*) for a contract. They are effectively *making a market*, since these quotes make transactions possible. In OTC markets, dealers are the primary liquidity providers. The most sophisticated among such markets allow the dealers to quote prices,

¹ Ch. 13 in Harris (2003).

quantities, and other attributes continuously. For example, fixed-income products can be traded on Dealerweb or Bloomberg. In order markets or for highly bespoke products, the dealers quote on request, possibly one-sided only, for a specific quantity and with an expiration time. The quote, or the spread if the quote is two-sided, depends on the quantity. Similarly to speculators, dealers trade on their own behalf. Like speculators, they hold a portfolio (or an *inventory* of positions) and face the issue of trading counterparties that may be more informed than themselves. Unlike speculators, dealers are passive traders, in that they respond to their clients. Also, unlike speculators, dealers enjoy special regulatory status. Because dealers observe the demand flow of their clients, they are informed agents, often serving the needs of informed clients. The dealers' profit originates from the realized spread of their trade (which is usually lower than the quoted spread) but also from the specific information the dealer derives from the order flow. One specific type of flow originates from retail investors (who we introduce later in the chapter). These investors access the market indirectly through brokers. Brokers have special arrangements to direct market orders to dealers, who commit to executing them while offering certain price guarantees on the trade.

In summary, dealers are liquidity providers, and they are compensated for services through trading profits.

- *Brokers*² trade on behalf of their clients. When the broker receives an order from a client, together with information about the client's time and price preferences, it searches for the most effective channel to execute it in accordance with these preferences. For example, a client sends a broker an order to buy a certain number of shares of a company. The broker is a member of all major exchanges. It splits the large order into smaller orders and routes them to the various markets at times that meet the execution horizon of the client or its expected cost. Unlike dealers, brokers are intermediaries who take no risk by holding contract positions at any given time. The intermediation service they provide is beneficial; however, it comes with its own risks. First, brokers provide exchange access to non-member clients and they provide OTC dealer access to non-institutional clients. Institutional clients, too, may want to enlist a broker when interacting with dealers, since the broker

² Ch. 7 in Harris (2003).

anonymizes the clients. Further, broker intermediation solves bilateral settlement risk: money is exchanged for contracts after the trade occurs. Clients need to know and trust their counterparty to protect themselves from insolvency, reneging, or non-compliance. There is a small number of brokers compared to the number of traders, so that clients need to approve (and be approved by) only a few counterparties; a reduction in time, cost, and risk. This, of course, does not eliminate counterparty risk. It transfers it to the brokers. The brokers manage it by vetting the clients, and by requiring that clients deposit capital at the broker, which the broker uses in case of client insolvency. The brokers also *clear* and *settle* trades on behalf of the client. In addition to these services, brokers, and especially *prime brokers*, the subclass of brokers servicing hedge funds and other sophisticated investors, offer their clients other services:

- *Custodial services*. Brokers ensure receipt, recording, and safekeeping of securities.
- *Rehypothecation*. Clients may allow the brokers to use their securities for the brokers' own needs in exchange for fees or rebates. For example, brokers may use client securities as collateral for their own transaction or lend them to other clients.
- *Margin loans*. Brokers lend clients short-term capital to buy securities. They charge them SOFR³ plus a spread.
- *Location of short positions*. Clients may want to *short* stocks (i.e., sell shares first, and buy them back at a later time), with the expectation that future prices will be lower than current ones. Brokers enable these transactions by lending shares from a third party and making them available to clients. The clients then sell them in the open market, buy them back at a later time, and return them to the broker. After the initial sale, but before the buyback, the broker invests the cash proceeds at SOFR plus a spread. The client receives from the broker SOFR minus a spread.
- *Research reports and services*, as well as broker-specific data. These services used to be bundled in broker commissions but after the implementation of MIFID II regulation, they are now charged separately.

³ Secured Overnight Financing Rate (SOFR) is a measure of the interest rate charged for overnight cash loans.

- *Capital introductions*, in which brokers facilitate the connection between hedge funds and potential investors.

In summary, brokers offer diversified services, the most important of which is to facilitate clients' transactions. They are compensated by commissions, interest on cash balances, interest on lending, and *payment for order flow* (PFOF), which is the compensation brokers receive from market makers for routing orders to them.

Broker-dealers, also called *dual traders*, combine the previous two functions in a single entity. They act both on behalf of the client and on their own behalf. This introduces a tension. The dealer's arm is incentivized to use the broker's information in trading to its advantage. Maybe the simplest action is *front-running*: the dealer is aware of incoming buying or selling demand for a security, and buys it in advance before this demand manifests in the market and is reflected in prices. To mitigate this type of behavior, regulations are in place to safeguard the interest of the client. The most important law regulating brokers, dealers, and broker-dealers is the Securities Exchange Act of 1934 (or "1934 Act").

1.3.2 The Buy Side

The buy side usually trades with the sell side. You (the reader of this book) are likely to be a member of this group, even though certain dealers face quantitative challenges similar to yours. It is important to understand who the actors in the buy-side drama are, because you will continuously interact with them, and your excess returns will be the outcome of this interaction. We could classify the buy-side actors according to several criteria. For example, the sub-industry to which they belong: life insurers, mutual funds, hedge funds, and so on. I opt to classify them (subjectively!) based on the type of investing they perform.

- *Indexers* are passive investors. Their portfolios replicate the compositions of the benchmarks, or indices, generated by data providers like MSCI, S&P, Russell, CRSP, or from exchanges like FTSE 100, TOPIX, and Deutsche Börse. These indices are updated on a quarterly or biannual basis, and they comprise bond indices as well, like the Bloomberg Agg (until 2016 owned by Barclays). Several investment vehicles track indices; mutual funds and exchange-traded funds are the largest in terms of size. Large firms in this group are Blackrock, Vanguard, and State

Street. Indexers make up a large and growing share of the total asset base. According to estimates by Chinco and Sammon (2023), they represent over 37% of the U.S. stock market capitalization as of 2020.

- *Hedgers* are firms participating in markets with the primary objective of reducing financial risk originating from their core businesses. For example, currency risk is faced by any firm doing business internationally. Firms such as airlines and manufacturing companies purchase fossil fuels (gas, Brent, and West Texas Intermediate), whose price variability can be very disruptive. Hedgers primarily participate in derivative markets: futures, swaps, and options. Hedgers differ from other participants who also hedge, such as dealers or hedge funds, in that hedging is the primary activity they perform.
- *Institutional active managers* are firms investing on behalf of their clients. They run strategies that are sometimes benchmarked to commercial indices and hope to beat them. There is some evidence of underperformance of funds serving retail investors; see S&P SPIVA report⁴ or the Refinitiv study (Glow, 2023). Both show that over 60% of funds underperform their benchmarks over a one-year trailing basis. The out-performance of funds over one year is not persistent: as of January 2024, 91% of funds trail the performance of the S&P500 over the previous 15 years. On the other side, funds serving institutions seem to beat their benchmarks (Gerakos and Linnainmaa, 2021). The *tracking error* is a measure of the risk they can take when differing from their reference benchmarks. Otherwise stated, their portfolios can be expressed as the sum of the positions in a benchmark, and of discretionary positions of a “tracking portfolio”. A large tracking error gives the funds much discretion; a low one places their returns close to the indices, and makes them “index huggers,” or “closet indexers” (i.e., index funds in disguise).
- *Asset allocators* manage portfolios composed of securities in multiple asset classes. Within an asset class, the portfolio closely follows a representative benchmark. One can view asset allocators as managers of a portfolio of asset classes. The relative weight of these asset classes in the portfolio is either constant or changes slowly. Common asset classes are equities, bonds, commodities, and cash equivalents.⁵ In addition, asset allocators

⁴ <https://www.spglobal.com/spdji/en/research-insights/spiva/>.

⁵ *Cash equivalents* are highly liquid assets with low returns, like bank certificates of deposit, short-dated treasuries, or commercial paper.

invest in alternative asset managers like private equity firms, venture capital, hedge funds, and real estate.

- *Informed traders* include primarily hedge funds and principal trading firms. These firms are usually organized as partnerships, although a few are public companies. They face fewer constraints than institutional managers. Whereas principal trading firms only have general partners (GPs, the principals) investing their own money, hedge funds also have limited partners (LPs) who do not invest actively.⁶ These firms pursue absolute returns (i.e., not tracking a benchmark), which exhibit low correlation to the indices of major asset classes.⁷ Informed traders invest heavily in human capital, technology, and data to achieve this goal. They fulfill two major functions. The first one is *price discovery*. By using all information available to them, they generate estimates of the true value of securities. If the security prices differ from their estimates, they trade to exploit the mispricing. If the price is lower than their estimate, they buy the security. In the process, they increase its price and bring it closer to equilibrium. Mispricing can take many forms. If the same security is offered at different prices on different exchanges, *arbitrageurs* (a subset of informed traders) will try to exploit the difference; of course, this may not be easy to do, so the difference either persists, or disappears very quickly due to technology investment in low-latency trading. The second role of informed traders is *liquidity provisioning*. Supply and demand of certain assets is predictable to a certain degree. I provide examples in Section 1.4; hedge funds and market makers develop specialized strategies that predict imbalances, hold (or short) securities before the liquidity need materializes, and meet the liquidity needs at the event. The range of possible intervals between prediction and event can be vast—from sub-second for high-frequency market makers to weeks or months for hedge funds.
- *Retail investors* trade for their own account via retail brokers. In 2020, retail investors made up approximately 20% of total volume; the share

⁶ GPs in hedge funds are both principals (since they invest their own money) and agents (on behalf of the LPs). To resolve this conflict of interest, the SEC regulates the class of investors who can be LPs (high-net-worth, sophisticated investors), and gives LPs special privileges within the fund.

⁷ This is not always true in practice because (a) some hedge funds have an explicit market exposure and (b) some hedge funds have an asymmetric exposure to the market (Agarwal and Naik, 2004).

was slightly more than 10% in 2011. Several studies, across different national markets and periods, have shown that retail traders are consistently unprofitable (Barber and Odean, 2013); retail trader flow is uninformed. This is one of the reasons why it is highly sought after by dealers, who will pay the retail brokers for routing it to them (payment for order flow).

1.4 Where Do Excess Returns Come From?

Now that we have introduced the main actors in the play (usually a tragedy, rarely a comedy, and occasionally a farce) of investing, we can discuss the sources of excess returns. The “excess” qualifier means “in excess of portfolio invested in risk-free assets, such as short-dated U.S. treasuries.” This topic is central both to academic financial research and to practitioners. Academic finance is primarily concerned with the question of *efficiency*. In the words of Malkiel (1987):

A capital market is said to be efficient if it fully and correctly reveals all available information in determining security prices. Formally, the market is said to be efficient with respect to some information set, ϕ , if security prices would be unaffected by revealing that information to all participants. Moreover, efficiency with respect to an information set, ϕ , implies that it is impossible to make economic profits by trading on the basis of ϕ .

An exceptionally concise definition, if there ever was one. At its core is the “information set ϕ ”. This could be, for example, the set of all historical prices of the traded securities. Nowadays, this information can be obtained with *relatively*⁸ little effort. A different type of information set is publicly available information,⁹ defined in the United States as “any information that you reasonably believe is lawfully made available to the general public from: (i) federal, state, or local government records; (ii) widely distributed media; or (iii) disclosures to the general public that are required to be made by federal, state, or local law.” An even finer information set

⁸ The “relatively” here denotes a bit of sarcasm. Collecting long time series of good quality, accounting for corporate actions, is hard, expensive, and requires skill.

⁹ PRIVACY OF CONSUMER FINANCIAL INFORMATION UNDER TITLE V OF THE GRAMM-LEACH-BLILEY ACT. §163.3 (w) (1).

is the set of *all* information available to *any* investor. Academic research tries to determine the validity of the statement that “security prices would be unaffected by revealing that information to all participants.” Note that this does not mean that ϕ is not helpful to predict future prices. Indeed, there is empirical evidence that asset prices are predictable. However, the hypothesis is that current prices may not be affected. We do not trade in the direction of returns, up to the point that the investing opportunity disappears. This is unintuitive. Why would we not take advantage of an informative prediction? One reason is *risk*. Even if we have some information about the future return of an asset, the uncertainty around the prediction is too high for us to take advantage of it. For example, say that, to the best of our knowledge of ϕ , we expect the U.S. market to appreciate 8% next year, while our cash custodied at the broker will return a measly 2%. Does this imply that we will rebalance our portfolio to 100% a market-tracking asset like SPY? Hardly. The reason is that the standard deviation of market returns is 20%, a little too high for comfort. Risk, however, is not the only reason. Another one is *liquidity*. Indeed, the road to hell of an investor is littered with quite accurate predictions of assets that barely trade or do not trade at all. A famous example is the spin-off of Palm (a now-defunct mobile device company) by 3Com (a telecom equipment maker, also defunct) in 2000. 3Com floated on the public market 5% of the shares of Palm, while retaining the other 95%. Right after the initial public offering (IPO), Palm had a market value of \$54B, while 3Com had a market capitalization of \$28B. The implied value of 3Com assets was –\$22B, even though the company had no debt, \$1B in cash, and positive cash flow. Either 3Com was dramatically undervalued, or Palm was dramatically overvalued. An investor could have therefore bought 3Com shares and shorted Palm for an equal amount. The portfolio comprising these two assets was a synthetic asset whose return could be predicted. There was a problem, however. Palm shares were in short (pun intended) supply. In order to short a share, the investor must first borrow it, at a rate decided by the lender. If quoted at all, these rates were so high as to make the trade either unattractive or impossible.

Risk and liquidity are not the only two factors limiting the exploitation of information. We list three more. The first one is *funding*. Consider a scenario in which certain assets, or certain portfolios,¹⁰ have lost much

¹⁰ As in the case of the pair 3Com–Palm, we can interpret portfolios as synthetic assets.

of their value due to market distress. We are managing a small hedge fund, which has also lost money in this environment. Based on historical examples, we have a strong belief that such assets will rebound. Such scenarios occur quite regularly, especially in “deleveraging spirals”. However, we do not have much capital available to post as margin. In addition, we need a capital buffer in order to withstand a possible additional loss in the very short term. Funding constraints prevent us from buying the asset, in spite of our accurate forecast.

A significant source of excess returns arises from *flow predictability*. Some agents, notably institutional investors and market makers, but not only these, will trade known securities on known dates. Speculators can then take advantage of this information by providing liquidity beforehand.¹¹ One of the most important instances of this is *index rebalancing*. Several index providers update the weights of their indices on predetermined dates using well-defined rules. Some securities are added to the index, others are removed, and finally most of the remaining ones have an updated weight. The term used for this process is *index reconstitution*. For example, TSLA was added to the NASDAQ 100, effective July 15, 2013. The announcement was made on July 10, 2013; but several investors could have forecasted the event well before that date. These investors would then purchase TSLA shares and sell them at the closing auction of July 15, 2013. The ETF, mutual funds, and bespoke products that track the index have an obligation to buy TSLA on the close of that day, and the resulting demand is likely to push up the stock price. The informed investors providing liquidity do not do so risk-free. They hold the stocks until the effective date, and over these days are exposed to the risk that TSLA may suffer from company-specific or industry-specific losses. Moreover, there is the remote risk that the reconstitution be cancelled or postponed. The size of passive investing is large and its estimate ranges from 17.5% (Novick, 2017) to 38% (Chinco and Sammon, 2023) of total assets under management. The buyer of index products bears the indirect cost of such rebalancing (Li, 2021). This is just one prominent example of predictable flows, but several others exist, usually smaller in size, but also not as widely known as index rebalancing. Their common feature is the existence of

¹¹ Note that this is not the same as *front running*. The latter also consists of buying or selling a security based on a demand forecast of said security; but in the latter case, the forecast relies on *non-public* information.

institutional or procedural constraints (sometimes driven by internal processes, other times by regulatory requirements) that introduce predictability in the demand of securities.

Finally, we consider a last source of excess returns: *informational advantage*. This means that the investors not only differ by their risk attitude, their tolerance to illiquidity, and their funding level, but also by their information sets, and some investors have better, more accurate forecasts of future prices or returns. This is what is often meant by “statistical arbitrage,” the ability to predict returns accurately based on insights our competitors do not have.

In summary, even assuming accurate information owned by participants, some of them cannot exploit this information. We have listed several possible causes of return predictability:

- pure arbitrage;
- heterogeneous risk preferences;
- liquidity;
- funding;
- flow predictability;
- informational advantages.

These categories are not exclusive. For example, flow predictability and liquidity are related, albeit not identical; and the distinction between being compensated for risk-taking and holding an actual information advantage is usually unclear. However, these broad classes can help us reason about one strategy’s edge.

1.5 The Elements of Quantitative Investing

The investment process is usually viewed as a highly structured process. There are separate components, the development of which is the responsibility of separate teams. In Figure 1.1, I show a possible organization of these components, which I follow in the organization of the book. I review them below.

- *Data*. The essential inputs to investing are under the “Data” section to the extreme left.




Figure 1.1 The components of the investment process.

- *Prices and trading volumes* are often collected at regular intervals (e.g., every minute, every 5, 10, 15 minutes, or daily). For high-frequency strategies it may be necessary to use order-level exchange data.
- *Characteristics* are numerical vectors associated with a security and a timestamp. Consider them as descriptors of the security. For a stock, a characteristic may be a measure of “quality”, like the free cash flow generated by the firm in the most recent quarter, divided by the market capitalization of the firm. Another widely used characteristic is the realized return of the stock over a certain interval, for example the past six months.
- *Time series* differ from characteristics (which are multivariate time series) in that they are not associated with individual securities, but rather provide additional information entering the investment process. Examples of time series are the Consumer Price Index (CPI), which can be used to estimate the inflation rate in the United States; the yield of the 10-year Treasury bond; the Federal Funds Effective Rate, the overnight rate for unsecured lending of reserves among commercial banks; and the VIX, a forward-looking measure of U.S. equities market volatility.
- *Unstructured data* are the dark matter of financial data. Prices, characteristics, and times series are structured data: numerical, categorical, or ordinal (i.e., rankings), and in tabular form. Unstructured data are usually character sequences representing natural language (e.g., earnings transcripts and firm news) or images (e.g., satellite images), video/audio files, or multimodal data (i.e., a combination of all of the above formats).

- *Before the trade.* Data are used to develop three components that enter the portfolio construction during the trade. Because they precede the trading process, I classify their development as being “before the trade”.
 - *Risk.* The word “risk” can mean many things. In the context of this book, we will use portfolio volatility as a proxy for risk. Estimating this volatility for an arbitrary portfolio is a challenging task.
 - *Expected returns.* In order to be profitable, a trading strategy needs to have informative predictions of future returns. This is often viewed as the paramount concern of a quantitative investor, and to a large extent it is. In nearly all modern trading systems there is more than one estimate of expected returns. The number of estimates can run into the thousands or millions.
 - *Transaction costs and market impact.* Trading securities is expensive, and these costs are unavoidable when deploying a strategy. Among other things, transaction costs determine whether a predicted return at some horizon can be turned into a profitable strategy or not; and what is the maximum profit that can be extracted from such a strategy.
- *During the trade.* The three components developed in the previous stage are combined in the portfolio construction phase. This happens “during the trade”, because the portfolio construction procedure results in real-time trading decisions, and these decisions determine the Profit and Loss (PnL) of the strategy. The decisions taken in the portfolio construction process are:
 - *Incorporation of risk constraints.* A strategy’s PnL is a function of the maximum risk that it can take. This is usually represented in the portfolio construction problem either in the form of constraints, or in the form of penalties added to the objective function. Risk constraints and penalties can have a very material impact on the performance of the strategy.
 - *Signal aggregation.* We use the term *signal* for a model of expected returns. As I mentioned above, there can be many signals for the returns of a single asset. A problem encountered in practice is combining such signals into a single signal.
 - *Hedging decisions.* Certain trading strategies have *exposure* to systematic risk. In layman’s terms: they can lose money because their returns are correlated to market-wide sources of risk. Some of these sources can

be *hedged*, which means that such risk can be counterbalanced. This is an important concept in portfolio construction.

- *After the trade.* The trading process generates a time series of PnL, both for the overall portfolio and for its constituents.
 - *Performance attribution.* In the portfolio construction phase we estimate expected returns. This is an *ex-ante* exercise. *Ex post*, we observe the actual PnL of the portfolio. Performance attribution is the practice of tracing back performance to its possible sources, to see what worked and what did not, so that we can learn from the experience. Moreover, performance attribution can also be employed to assess whether we have skill in sizing our bets.
 - *Intertemporal volatility allocation and leverage.* How should we allocate risk across periods? This decision is very consequential to capital growth. A closely related question is that of leverage, defined as the ratio between Gross Market Value (GMV) of the portfolio and Assets under Management (AuM).

I stress that different arrangements are possible; for example, Narang (2024) employs a simpler scheme, and Pedersen (2015) has yet another one. Some of the individual elements I introduce are present in these models of quantitative investing.

The Takeaways

1. Market participants can be broadly classified into sell-side and buy-side participants. The *sell-side participants* are

- *Dealers*
- *Brokers*
- *Broker-dealers*

The *buy-side participants* are

- *Indexers*
- *Hedgers*
- *Institutional active managers*
- *Asset allocators*
- *Informed traders (e.g., hedge funds, principal trading firms)*
- *Retail investors*

2. *Excess returns* arise from five major sources:

- (a) *Risk*
- (b) *Liquidity*
- (c) *Funding constraints*
- (d) *Predictable flows*
- (e) *Informational advantage* in predicting future returns

3. *Quantitative investing* employs analytical models that can be partitioned into three broad groups:

- (a) *Risk measurement* models and data
- (b) *Market impact* models
- (c) *Models of expected returns*

4. The *investment process* has several stages:

- (a) *Data collection*: Prices, trading volumes, asset characteristics, time series data, and unstructured data
- (b) *Pre-trade analysis*: Risk assessment, expected returns estimation, and transaction cost analysis
- (c) *Portfolio construction*: Aggregation of signals, application of risk constraints, and hedging
- (d) *Post-trade evaluation*: Performance attribution and risk adjustments based on intertemporal volatility and leverage considerations

Chapter 2

Univariate Returns

The Questions

1. What are the definitions and types of returns, including dividend-adjusted and excess returns?
2. What are the stylized facts of stock returns, and why are they important?
3. How do we estimate prices and returns while accounting for market microstructure effects?
4. What are Conditional Heteroskedastic Models (CHMs) and how do they model volatility?
5. How does the GARCH(1, 1) model capture the stylized facts of returns?
6. What is realized volatility and how is it estimated using high-frequency data?
7. How can state-space models and the Kalman filter be used for variance estimation?
8. How does GARCH(1, 1) relate to Exponentially Weighted Moving Averages (EWMA) in volatility estimation?

We begin with models of univariate returns for two reasons. First, single-asset returns are the basic constituents of portfolios. We cannot hope to understand portfolio behavior without a solid understanding of their building blocks. Therefore, it is necessary to summarize the salient empirical properties of stock returns and the most common processes employed to model them, specifically to model volatility effectively. Second, these models have general applicability and are even more useful when combined with other families of models for multivariate returns. GARCH and exponential moving averages are essential tools for the working modeler. In the process, I introduce models that justify their use. Exponential moving averages find their motivation in linear state-space models, while GARCH is an instance of a non-linear state-space model. These models will be your friends for life. The chapter is organized into four parts. First, we lay out definitions of returns. Second, we summarize some “stylized facts” (empirical features of returns that are ubiquitous and relevant to risk management). As part of basic volatility modeling, we will introduce GARCH models and realized volatility models. Given that these topics are well covered in textbooks, the goal here is to present the essentials, their associated insights, and provide a jumping-off point for the reader. Finally, I will touch on the state-space models for variance estimation.

2.1 Returns

2.1.1 Definitions

We have a set of n assets and a currency, also called the numeraire.¹ Throughout this book, we use dollars as the base currency. It is customary to assume that each of these assets is infinitely divisible. We buy today the equivalent of a unit of currency for asset i . We denote the value of the asset tomorrow (or some future period) $1 + r_i$. An equivalent way to

¹ This word comes to English from the Latin *numerarius*, or “a number”, “a unit”, through the French *numéraire*.

define returns is from the closing prices of security i on days 0 and 1, which we denote $P_i(0)$ and $P_i(1)$, respectively. The *return*² is defined as

$$r_i(1) := \frac{P_i(1) - P_i(0)}{P_i(0)}$$

We extend this definition to the case in which the security pays a dividend. The holder of the asset receives an amount $D_i(1)$. The *dividend-adjusted return* is defined as

$$r_i(1) := \frac{P_i(1) + D_i(1) - P_i(0)}{P_i(0)}$$

In a universe of n assets, the vector of daily returns between times $t - 1$ and t is denoted $\mathbf{r}(t) := (r_1(t), \dots, r_n(t))$. For a portfolio $\mathbf{w} \in \mathbb{R}^n$, where w_i is a monetary amount invested in asset i , the PnL in a single period is given by the change in the value of the portfolio. The number of shares owned in asset i is calculated as $w_i/P_i(0)$. The value of the portfolio in period one is $\sum_i (w_i/P_i(0))P_i(1)$, and the change in value is

$$\sum_i \frac{w_i}{P_i(0)} P_i(1) - \sum_i w_i$$

In vector form, this equals $\mathbf{w}^T \mathbf{r}$. The *volatility* of a random return is its standard deviation: $\sigma := \sqrt{E[(r - E(r))^2]}$. The *variance* is the square of the volatility. Occasionally, when the approximation $E(r^2) \gg E^2(r)$ holds, we will approximate the volatility by the second moment of the return: $\sigma^2 \simeq E(r^2)$. The *compounded return* is defined as the value of one unit of numeraire after T period, had it been invested in a security yielding returns $r(t)$, i.e.,

$$\begin{aligned} r(1 : T) &:= \frac{P(T)}{P(0)} - 1 \\ &= \frac{P(T)}{P(T-1)} \frac{P(T-1)}{P(T-2)} \cdots \frac{P(1)}{P(0)} - 1 \\ &= \prod_{t=1}^T (1 + r(t)) - 1 \end{aligned}$$

² Definitions of returns, log returns, and dividend-adjusted returns are in Connor et al. (2010) and Ruppert and Matteson (2015).

Index

R^2 , *See* Coefficient of
Determination

A

ACF *See* Autocorrelation
Active Managers 10
ADR 166, 181
Aggregational Gaussianity 28,
32, 52
Akaike Information Criterion
(AIC) 105
Alpha 62, 64, 80
Intraday 335
Orthogonal 64, 69, 236, 302
Spanned 64, 68, 69, 71, 230,
236, 302
Alternate Trading System 279
American Depositary Receipt
See ADR
Amihud's Liquidity Measure
143
Announcement Date 193
Arbitrageurs 11
Asset Allocators 10
Assets Under Management 18
AUM *See* Assets Under
Management
Autocorrelation 27

Absolute univariate returns
27
Univariate Returns 26

B

Backtesting
Cross-Validation 200, 219
Protocol 195
Rademacher Anti-Serum
(against backtesting bites)
200
Walk-Forward 195, 200, 219
Walk-Foward 198
Bayesian Information Criterion
105
Beginner's Mind, xix
Benchmarks 9
Beta 77
Betas
Adjusted-Dollar 343
Bias Statistic 96
Bid-Ask Spread 24
Bloomberg 7
Bonds 3
Book-to-Price Ratio 235
Borrow Costs 194
Broker-Dealers 9
Brokers 7

Brownian Process 320
Buy Side 6, 9

C

CA *See* Closing Auction
Cantelli's Inequality 56
Capital Asset Pricing Model 56
CAPM *See* Capital Asset
Pricing Model
Cash Equivalents 10
CDS *See* Credit Default
Swaps
Central Limit Theorem 30, 185
Characteristics 16
Child Order 278
CHM *See* Conditional
Heteroscedastic Models
Clearing 8
Closet Indexers 10
Closing Auction 112
Clustering
K-Means 174
Coefficient of Determination
104, 240
Condition Number of a Matrix
268
Conditional Heteroscedastic
Models 30

Heteroskedastic noise 85
 HFT *See* High-Frequency
 Trader
 HIC SVNT LEONES 135, 321
 High Watermark 327
 High-Frequency Trader 335
 History 189

I

IC *See* Information Coefficient
 Idempotent Operator 74
 Idiosyncratic Variance 70, 77,
 79
 iid rv, xxiv
 Index
 Rebalancing 14
 Reconstitution 14
 Index Huggers 10
 Indexers 9
 Indices 9
 Information Coefficient 193,
 198, 200, 204, 238
 Information Ratio 54, 55, 240,
 348, 353
 Empirical 352
 Information Set 12
 Interest Rate Swaps 4
 IR *See* Information Ratio
 IRS *See* Interest Rate Swaps

K

Kalman Filter 26, 42, 46
 Optimal Gain 48
 Kelly criterion 223, 313,
 332
 Kolmogorov-Smirnov distance
 34

L

Lagrange Multiplier 225
 Leverage 18
 Leverage Effect 28
 Likelihood 83
 Limit-Order Book 5, 24, 112
 Limited Partners 11
 Linear Regression 82, 232
 As Projection 84
 Decomposition 86
 Features 83
 Frisch-Waugh-Lovell Theorem
See Frisch-Waugh-Lovell
 Theorem
 Random Design 115
 Linear State-Space Models 21,
 42, 44
 Liquidity 4, 13

Liquidity Provisioning 11
 Liquidity Spirals 143
 Loadings
 Definition 63
 Orthonormal 72
 Z-scored 72
 Z-scoring 72
 LOB *See* Limit-Order Book

M

Mahalanobis distance 101
 Market Efficiency 12
 Market Impact 17, 277
 Market Makers 6
 Market orders 7
 Marketable Order 279
 Massart's lemma 205
 Maximum Likelihood Estimator
 34, 38
 Meta-Order *See* Parent Order
 MIFID II 8
 MLE *See* Maximum Likelihood
 Estimator
 Moore-Penrose Pseudo-Inverse,
 xxiii 119
 MSE Statistic 96, 97

N

Net Market Value 23
 Net Market Value (NMV) 67
 Newey-West Estimator 125,
 135
 NMV *See* Net Market Value
 Norm
 p -Norm, xxiv
 Frobenius 149
 Operator, xxiv 268
 Unitarily Invariant 149
 Numeraire 21, 67

O

Optimization
 Mean-Variance 222,
 224–227, 231, 236, 237,
 242, 244–246, 251, 254,
 263, 324
 Orthogonal Complement 86
 OTC *See* Over-the-Counter
 Over-the-Counter 5, 279

P

Parent Order 278
 Partial Correlation 229
 Participation Rate 280
 Passive Investing 14
 Payment for Order Flow 9, 12

PCA *See* Principal Components
 Analysis
 K-Means 174
 Probabilistic 152
 Penalty
 1-Norm 265
 Norm 264
 Squared Norm 262
 Performance Attribution 61,
 68, 72, 75, 93, 335
 Characteristics 335
 Maximal 340, 343
 Paradoxes 336
 Selection-Sizing 347
 Time Series 335
 Permanent Market Impact 278
 PFOF *See* Payment for Order
 Flow
 PnL 17, 18, 22
 Factor 68, 76
 Idiosyncratic 76
 Position 335
 Trading 335
 Portfolio 22
 basis 97
 Characteristic 235
 Construction 61, 64
 Eigenportfolios 97
 Factor-Mimicking 103, 104,
 109, 116, 121, 229–231,
 233, 234, 236, 237, 245,
 301, 336
 From Sorts 235
 Market 103
 Minimum-Variance 107–109
 Penalties 230, 261
 Production 97
 Thematic 103
 Portfolio Management 75, 80,
 220, 241, 250, 257
 POV *See* Participation Rate
 Precision Matrix 80, 225
 Price
 Ask 6
 Bid 6
 Price Discovery 11, 335
 Price Manipulation 281
 Principal Component 150, 151
 Principal Component Analysis
 149, 150, 185
 Principal Components Analysis
 152, 153, 155, 156, 158, 169,
 171–174, 177, 179
 Principal Trading Firms 11
 Probabilistic PCA 154, 179