
Frank J. Fabozzi, CFA
Yale University and *Journal of Portfolio Management*

Sergio M. Focardi
The Intertek Group

Petter N. Kolm
Yale University

Trends in Quantitative Finance



Statement of Purpose

The Research Foundation of CFA Institute is a not-for-profit organization established to promote the development and dissemination of relevant research for investment practitioners worldwide.

The Research Foundation of CFA Institute and the Research Foundation logo are trademarks owned by The Research Foundation of CFA Institute. CFA®, Chartered Financial Analyst®, AIMR-PPS®, and GIPS® are just a few of the trademarks owned by CFA Institute. To view a list of CFA Institute trademarks and a Guide for the Use of CFA Institute Marks, please visit our website at www.cfainstitute.org.

© 2006 The Research Foundation of CFA Institute

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, without the prior written permission of the copyright holder.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

ISBN 978-1-932495-49-2

Printed in the United States of America

21 April 2006

Editorial Staff

Elizabeth A. Collins
Book Editor

David L. Hess
Assistant Editor

Kara H. Morris
Production Manager

Lois Carrier
Production Specialist

Biographies

Frank J. Fabozzi, CFA, CPA, is adjunct professor of finance and Becton Fellow in the School of Management at Yale University and editor of the *Journal of Portfolio Management*. Prior to joining the Yale faculty, he was a visiting professor of finance in the Sloan School at the Massachusetts Institute of Technology. Professor Fabozzi is on the advisory council for the Department of Operations Research and Financial Engineering at Princeton University and has authored and edited numerous books in finance. In 2002, he was inducted into the Fixed Income Analysts Society's Hall of Fame. Professor Fabozzi holds a doctorate in economics from the City University of New York.

Sergio M. Focardi is a founding partner of the Paris-based consulting firm The Intertek Group. Mr. Focardi lectures at the Center for Interdisciplinary Research in Economics and Finance at the University of Genoa (Italy) and is a member of the editorial board of the *Journal of Portfolio Management*. He has written numerous articles on econophysics and is the coauthor of *Modeling the Markets: New Theories and Techniques*; *Risk Management: Framework, Methods and Practice*; *The Mathematics of Financial Modeling and Investment Management*; and *Financial Modeling of the Equity Markets: From CAPM to Cointegration*. Mr. Focardi holds a degree in electronic engineering from the University of Genoa and a postgraduate degree in communications from the Galileo Ferraris Electrotechnical Institute, Turin, Italy.

Petter N. Kolm is a doctoral student in finance at the School of Management, Yale University, and a financial consultant in New York City. Previously, he worked in the Quantitative Strategies group at Goldman Sachs Asset Management, where his responsibilities included research and development of new quantitative investment strategies for the group's hedge fund. Dr. Kolm is a member of the editorial board of the *Journal of Portfolio Management* and is a coauthor of *Financial Modeling of the Equity Markets: From CAPM to Cointegration*. He holds a doctorate in mathematics from Yale University, an MPhil in applied mathematics from the Royal Institute of Technology in Stockholm, and an MS in mathematics from ETH Zurich.

Contents

Foreword	v
Introduction	vii
Chapter 1. Forecasting Financial Markets	1
Chapter 2. General Equilibrium Theories: Concepts and Applicability	13
Chapter 3. Extended Framework for Applying Modern Portfolio Theory	21
Chapter 4. New Territory: Modeling for Portfolio and Tactical Asset Management	35
Chapter 5. New Territory: Forecastability and the Long-Term Perspective	44
Chapter 6. Machine Learning	57
Chapter 7. Model Selection, Data Snooping, Overfitting, and Model Risk	67
Chapter 8. Predictive Models of Return	76
Chapter 9. Model Estimation	91
Chapter 10. Practical Considerations When Using Optimization Software	104
Chapter 11. Models in Practice: Industry Survey Results	112
Chapter 12. Quantitative Modeling Today and Tomorrow	119
References	124



This publication qualifies for 5 PD credits under the guidelines of the CFA Institute Professional Development Program.

Foreword

Mark Kritzman, CFA, my predecessor as research director at the Research Foundation of CFA Institute, is a wit as well as an intellect and once referred to those casually acquainted with quantitative finance as “dilequants” (rhymes with dilettantes). If you dabble in quantitative methods and wonder whether you might be so characterized, you shouldn’t be insulted. I’m a dilequant too. Understanding and applying quantitative techniques in finance takes a lifetime of study and mastery, and most practitioners would do well to strive for understanding what quantitative methods in finance are, and what they are best used for, rather than trying to achieve this mastery on their own.

But for quantitative methods to be used and appreciated in the investment community, one needs a primer on the topic for a nontechnical audience. The current monograph achieves this difficult goal. Its authors, Frank J. Fabozzi, CFA, Sergio M. Focardi, and Petter N. Kolm, have translated the often highly technical jargon and mathematical language used by “quants” into plain English.

Quantitative finance is broadly applied in three areas: (1) screening universes of securities to help select those one wants to buy (or sell short) in an effort to add alpha relative to a benchmark, (2) portfolio construction, in which optimization and related methods are used to build efficient portfolios of those securities, and (3) pricing derivatives.

The current monograph focuses, strongly but not exclusively, on portfolio construction. Fabozzi, Focardi, and Kolm pay considerable attention to optimization in the presence of estimation error, a topic raised most visibly by Richard Michaud in his January/February 1989 *Financial Analysts Journal* article, “The Markowitz Optimization Enigma: Is ‘Optimized’ Optimal?” Approaching the problem from a different angle, Fischer Black and Robert Litterman, in their September/October 1992 *Financial Analysts Journal* article “Global Portfolio Optimization,” also addressed the issue of estimation uncertainty in portfolio construction, as did J. David Jobson and Bob Korkie in a series of articles in the early 1980s. Fabozzi, Focardi, and Kolm expand on all of these concerns. And increased interest in alternative assets, such as hedge funds, for which the standard assumption of a normal distribution of returns may not apply, creates a need for “robust” optimization methods, to which the authors of this monograph devote considerable attention.

Another topic addressed by Fabozzi, Focardi, and Kolm is the use of advanced econometric techniques to try to add alpha by forecasting security (or asset-class) returns. Although the standard assumptions of portfolio theory—the efficient market hypothesis, the no-arbitrage condition, and general equilibrium models of asset pricing, such as the capital asset pricing model—posit a world in which returns are not forecastable, these assumptions do not always hold up. Practitioners have

made money by forecasting. The current monograph provides a primer on some of the more widely used forecasting techniques by covering such important issues as model selection, biases in models, and data mining and snooping.

Finally, in an innovative section, the authors provide results of a survey in which investment management organizations reveal what quantitative techniques they use and what challenges they face in using them.

In summary, Fabozzi, Focardi, and Kolm provide an excellent and comprehensive survey of the challenges one meets in using quantitative methods for portfolio construction and forecasting. By covering a wide variety of methods rather than advocating a particular one, the monograph reflects an inclusive and thoughtful approach.

The Research Foundation is very pleased to present *Trends in Quantitative Finance*.

Laurence B. Siegel
Research Director
The Research Foundation of CFA Institute

Introduction

The aim of this monograph is to introduce practitioners to recent developments in the modeling of equity returns for the purpose of asset management. We have tried to provide a plain-English, formula-free review of quantitative methods without sacrificing conceptual rigor. In addition to discussing methodology, the monograph includes the results of an *ad hoc* survey taken in the first half of 2005 of equity modeling at 21 large asset management firms in the United States and Europe.

As a profession, asset managers have traditionally tried to “beat the market”—that is, to earn returns in excess of returns obtained by an indexed strategy. Their ability (real or supposed) to construct portfolios that earn excess returns is the reason investors entrust assets to them and the justification for active management fees in excess of index fund fees. This effort to outperform the market is one reason for the growing use of modeling techniques in asset management.

Market Efficiency and Unpredictability

Under the assumption that modern financial markets are efficient, mainstream finance theory has traditionally held that markets cannot be beaten. Although excess returns might indeed be achieved, they are considered to be, on average, proportional to risk: Markets embed a risk–return trade-off in which investors demand, and markets supply, excess expected returns for taking risk. In an efficient market, the risk–return trade-off also implies that above-market returns cannot be achieved without taking additional risk.

The assumption of market efficiency is associated with the notion of the *unpredictability* of financial markets. Mainstream theory maintains that markets must be unpredictable because if markets were predictable, they could not be efficient and returns in excess of market returns could be made without taking additional risk.

The notion of market efficiency has given rise to “passive” asset management strategies because if markets are efficient, trying to beat the market is futile. Better to reduce management costs by investing so as to match the performance of broad indices.

Given the double-digit returns that the U.S. equity markets were providing in the last two decades of the 20th century, few practitioners were overly concerned about the debate on market efficiency. As the 1990s drew to a close, however, the academic view of market efficiency and market predictability began to change under the weight of empirical evidence and fresh theoretical insights. The market downturn after 2000 forced asset management firms to reevaluate their investment processes in an effort to reduce costs and produce returns in unrewarding markets. Because quantitative methods can help in both tasks, many firms began to take a closer look at these methodologies.

Actually, a complete conceptual overhaul of our thinking about equity price processes is needed. The practice of investment management has to be reconciled with a new theoretical concept of asset returns—namely, that the trade-off between risk and return is dynamic and does not exclude the possibility that asset returns are, to some extent, forecastable. This monograph provides an overview of the recent changes in finance theory and the modeling techniques that the industry is using or beginning to experiment with in an attempt to capture the limited forecastability in financial markets.

New Concepts of Risk and Return

The only really general observation we can make about market efficiency is the absence of arbitrage—that is, in the financial markets, one cannot make a sure profit with no net investment. There is no free lunch. Pragmatically, therefore, whatever strategy investors adopt, they always face the possibility of losing money. Although finance theory states that investors (or asset managers) cannot beat the market without risk (because doing so would entail arbitrage), it does admit that an investor can beat the market, on average, by taking risk beyond the risk inherent in the market benchmark. Taking this additional risk means, of course, that the investor will suffer periods of underperformance as well as periods of superior performance relative to the market benchmark.

To make the critical decisions about how much risk to take, the profession clearly needs a quantitative framework for measuring risk and return—which is provided by probability theory and statistical techniques. The quantitative principles of investment under uncertainty were laid down by Markowitz (1952) more than 50 years ago; their adoption in full earnest requires the use of quantitative methods and modeling. But even today, the adoption of these techniques by the asset management community is patchy.

In a probabilistic quantitative framework, a number of concepts about markets have to be critically revised. Market efficiency does not imply that all investments are equivalent: Given one's risk–return preferences, some investments are preferable to others. Thus, we cannot state that all excess returns are equally offset by risk to the point where every investment has the same certainty-equivalent return. Some returns are less offset by risk than others. What remains true is that without investment and risk, one cannot make money. Equivalently, one cannot always, or even usually, beat the market.

To measure the ability of a manager to engineer a favorable risk–return trade-off, researchers introduced the concepts of *beta* (a measure of exposure to market risk) and *alpha* (a measure of return in excess of the market return, which can be interpreted as measuring skill in stock picking or asset allocation). All security or portfolio returns comprise a market part (beta) and a nonmarket part (alpha). The beta part of the return is caused by correlation with the relevant market benchmark

and thus arises from market exposure, not active management. The alpha part is the return “above and beyond” the beta part and represents the value added by the active manager.

Note explicitly that many realizations of asset management strategies will show positive alpha *ex post*. The key challenge of investment management, however, is to identify *ex ante* which strategies will produce positive alpha. Having generated a positive alpha *ex post* is not by itself a sign of a good active strategy: Such a result can sometimes be achieved simply by luck. Strategies can be considered alpha generators only if alphas are persistent. For this reason, performance measurement is a delicate issue. Because we cannot rely on always having access to long series of past performance, we try to gauge the true performance of an asset manager by correcting his *ex post* performance with an estimate of the risk associated with his strategy.

Models of equity returns can be static or dynamic. The models of standard finance are static; that is, the distributions of the model variables do not depend on the previous path of the same variables. A random walk is a typical example of a static model. Consequently, from the point of view of standard finance, alphas and betas are interpreted as static terms; they are constants that do not change over time.

However, we can also model the market with dynamic models. In these models, the variables do depend on their previous paths. If we use dynamic models, the concepts of alpha and beta have to be reinterpreted. In fact, in linear dynamic models, we typically find long-term equilibrium relationships plus short-term dynamics. The implication is that alpha and beta change over time. Moreover, if we add nonlinearities and higher statistical moments (such as skewness or kurtosis) or nonnormal distributions, we find that the risk–return trade-offs of assets cannot be described by the linear relationship implied by alpha and beta.

Dynamic models entail predictability of expected returns or of higher moments. This predictability is compatible with finance theory if it generates no arbitrage opportunities. And keep in mind that forecasting models do not necessarily offer better risk–return trade-offs than static models without predictability. True static alpha, if it exists, generates abnormal profits without the trading costs associated with dynamic strategies.

Generally speaking, given the large universe of investable stocks, capturing market opportunities requires optimization methodologies to fully exploit the risk–return trade-offs that modeling allows us to identify. Entrusting the management of large sums to automatic models and optimizers entails a high level of confidence in models, however, so the robustness of the quantitative models (that is, their relative insensitivity to a violation of one or more assumptions) has become an important concern for many firms.

A central theme of this monograph is the trade-offs that must be made among model complexity, model risk, and model performance. We return to this idea time and again—particularly in Chapter 3 on robust methods, Chapter 6 on machine learning, Chapter 7 on model selection, and Chapter 9 on model estimation.

Overview of the Monograph

The 12 chapters of this monograph develop the themes we have outlined. We begin by analyzing the concept of forecastability. We discuss the difficulty in predicting financial markets because the predictions themselves influence (modify) market behavior. This phenomenon, known as “self-referentiality,” does not mean that forecasting markets is impossible, only that there are constraints on the risk–return trade-offs offered by financial markets.

We argue that, counter to intuition, financial markets cannot be completely unpredictable yet at the same time contain a risk–return structure. If markets were totally unpredictable, for risk to be rewarded, they would have to exhibit different time-invariant expected returns. A static, immutable spread of returns between assets of different risk would lead to exponentially diverging prices and to exponentially diverging market capitalizations. This would occur whether stock returns provide alphas or not. Empirically, however, we do not find an exponential distribution of market capitalizations.¹ This observation leads to the conclusion that there is some forecastability in markets.

The idea that financial markets have some degree of forecastability has now gained broad acceptance. However, predictability is not automatically a source of profitability. We close Chapter 1 with a discussion of the need to carefully evaluate (1) the risk–return trade-off implied by the models and (2) transaction costs so as to ensure that strategies that look profitable on paper do not end up producing losses and/or inferior performance relative to a benchmark when applied in practice.

In Chapter 2, we outline the basic principles of general equilibrium theories. The objective is to improve understanding of the capital asset pricing model and the notion of market equilibrium. We then introduce the concept of the utility function, which represents the investor’s financial decision-making processes. The utility function has proved to be an important concept for the practice as well as the theory of finance. In fact, every optimization process depends on the specification of a utility function.

Despite their theoretical weight, *general equilibrium* theories are difficult to test and to use in practice. The reason is that the specification of the utility function remains abstract; it is an *a priori* assumption, one not based on empirical investigation. In the absence of an independent empirical evaluation of utility functions, general equilibrium is a theoretical framework that can always be used insofar as, in the absence of arbitrage, any price process can be rationalized as a general equilibrium.

¹Actually, we do find empirically that market capitalization follows a Pareto law. This Pareto law can be described intuitively by one of its properties: The size of an individual is inversely proportional to its rank. That is, the size of the second largest company is one-half the size of the biggest company, the third largest is one-third, and so on. Many phenomena, including economic phenomena, obey Pareto’s laws.

In Chapter 3, we describe the modern robust framework for Markowitz mean–variance optimization. We begin by describing the essentials of mean–variance optimization theory.

One critical aspect of the theory is the estimation of the variance–covariance matrix. Because in estimating the variance–covariance matrix the number of entries grows with the square of the number of assets that are candidates for portfolio inclusion, the matrix becomes rapidly unmanageable. We discuss robust estimation methods that allow one to reduce the number of independent covariance entries to be estimated.

A second critical component of the modern framework for Markowitz mean–variance estimation is *robust optimization*. Introduced recently in finance and still a subject of research, robust optimization places constraints on the results of the optimization process as a function of the uncertainty associated with parameter estimation. We discuss how robust estimation and robust optimization are two integrated aspects of robust methodologies.

In Chapter 4, we begin to explore models that detect forecastability in asset returns. We discuss the types of delayed responses that markets can exhibit to past values of variables, such as prices or returns. Forecastability is thus exploited by strategies based on momentum, reversal, co-integration, and mean reversion. We then discuss the issue of model complexity and sample size—that is, the size of the available historical dataset. There is a relationship between the size of the sample used for estimation and the complexity of the models we can estimate. If the sample is large, we can estimate a complex model; otherwise, we can estimate only the essentials.

In Chapter 5, we review issues related to modeling at different time horizons. Most models currently in use are estimated and reestimated on moving “windows” of historical data. We discuss the conditions that allow the estimation of slowly changing models and models that exhibit sudden regime shifts. Then, we discuss the behavior of stock markets at long time horizons and the concept of time diversification (i.e., the concept that financial risk is statistically smaller in the long run than in the short run because the ups and downs tend to offset each other, on average).

In Chapter 6, we provide an overview of *machine learning* and its applications in finance. Machine learning is a universal modeling strategy that does not depend on any domain-specific theory. Therefore, when applied to finance, the models do not use finance theory but rely on purely statistical analysis of financial phenomena. Machine-learning methods place constraints on model complexity to ensure that they retain some forecasting capability. We discuss a number of specific techniques, including neural networks, decision trees, clustering, genetic algorithms, and support vector machines. We also provide a perspective on artificial intelligence and techniques for handling unstructured (e.g., textual) data and text-related technologies.

In Chapter 7, we review the process of model selection and its pitfalls. We discuss how to deal with *data snooping* and avoid *survivor biases*. We also cover risk mitigation in modeling and, extending the discussion begun in Chapter 4, consider model complexity and the size of sample data.

Chapter 8 offers an overview of models used in equity return forecasting. Among the families of models discussed are the widely used models that regress returns on predictors and models that exploit momentum and reversal phenomena. We also discuss complex models that, although not widely used in asset management today, are beginning to make their way into practice. Among these are autoregressive models, factor models, hidden-variable models, and regime-switching models.

Model estimation is the subject of Chapter 9. Although (in keeping with the nature of this monograph) this chapter does not contain formulas, it does provide an overview of the concept of estimation and of the sampling distribution. We then present the most widely used estimation methods: the least-squares method, the maximum-likelihood estimation method, and the Bayesian estimation method. The chapter closes with a description of the estimation of regressions and other related models introduced in previous chapters.

Optimization, and in particular robust optimization, is becoming an important component in portfolio management applications. Chapter 10 presents the conceptual framework of optimization and gives practical suggestions for implementation and software selection. The development of robust methods for estimation and optimization is one of the major achievements of modern financial modeling. Robust technologies assume that models and the inputs themselves (like humans) are uncertain; they evaluate the consequences of errors in the models and introduce corrections that mitigate the potentially negative effects of model and estimation errors.

One of the objectives of this monograph is to provide a reading of how quantitative methods are making their way into the investment management process. Chapter 11 presents the results of an *ad hoc* market survey covering the use of quantitative methods in three areas: equity return forecasting, model risk mitigation, and optimization. Twenty-one asset management firms in the United States and Europe shared information on what modeling approaches they are actually using and experimenting with. Survey results are discussed and summarized in a table.

Finally, Chapter 12 considers the state of quantitative modeling today, with a discussion of modeling for portfolio management and for the profession in general, and suggests some possible future developments.

Acknowledgments

We wish to express our appreciation to the Research Foundation of CFA Institute for funding this project and providing us with guidance in preparing the manuscript. Mark Kritzman, CFA, then chair of the Research Foundation, first suggested this project to us after reading some of our earlier research and reviewed several of the earlier chapters. Laurence Siegel who replaced him as chair of the Research Foundation, made helpful suggestions on the first draft of the manuscript. We also received insightful comments from Dennis McLeavey, CFA, of CFA Institute and Peter Bernstein.