

The Portable  
**Financial Analyst**

*What Practitioners  
Need to Know*

Second Edition

MARK P. KRITZMAN



John Wiley & Sons, Inc.



The Portable

# Financial Analyst

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, please visit our Web site at [www.WileyFinance.com](http://www.WileyFinance.com).

The Portable  
**Financial Analyst**

*What Practitioners  
Need to Know*

Second Edition

MARK P. KRITZMAN



John Wiley & Sons, Inc.

Copyright © 2003 by Mark P. Kritzman. 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](http://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, e-mail: [permcoordinator@wiley.com](mailto:permcoordinator@wiley.com).

**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. Neither the publisher nor author 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 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 books.

For more information about Wiley products, visit our web site at [www.wiley.com](http://www.wiley.com).

*Library of Congress Cataloging-in-Publication Data:*

Kritzman, Mark P.

The portable financial analyst: what practitioners need to know / Mark Kritzman.—2nd ed.

p. cm.—(Wiley finance series)

Includes bibliographical references and index.

ISBN 0-471-26760-0 (cloth)

1. Investment analysis. 2. Portfolio management. 3. Investments.

I. Title. II. Series.

HG4529.K747 2003

332.6—dc21

2003057656

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

# Contents

<b>Foreword</b>		
<b>Time and Magic: What This Book Is About</b>		<b>vii</b>
<b>Preface</b>	<b>xi</b>	
<b>CHAPTER 1 The Nobel Prize</b>	<b>1</b>	
<b>CHAPTER 2 Uncertainty</b>	<b>7</b>	
<b>CHAPTER 3 Utility</b>	<b>17</b>	
<b>CHAPTER 4 Lognormality</b>	<b>27</b>	
<b>CHAPTER 5 Return and Risk</b>	<b>33</b>	
<b>CHAPTER 6 Higher Moments</b>	<b>41</b>	
<b>CHAPTER 7 Duration and Convexity</b>	<b>49</b>	
<b>CHAPTER 8 The Term Structure of Interest Rates</b>	<b>57</b>	
<b>CHAPTER 9 Serial Dependence</b>	<b>66</b>	
<b>CHAPTER 10 Time Diversification</b>	<b>75</b>	
<b>CHAPTER 11 Regressions</b>	<b>85</b>	
<b>CHAPTER 12 Factor Methods</b>	<b>94</b>	

---

<b>CHAPTER 13 Estimating Volatility: Part I</b>	<b>102</b>
<b>CHAPTER 14 Estimating Volatility: Part II</b>	<b>110</b>
<b>CHAPTER 15 Hypothesis Testing</b>	<b>114</b>
<b>CHAPTER 16 Future Value and the Risk of Loss</b>	<b>123</b>
<b>CHAPTER 17 Event Studies</b>	<b>132</b>
<b>CHAPTER 18 Simulation</b>	<b>139</b>
<b>CHAPTER 19 Value at Risk</b>	<b>148</b>
<b>CHAPTER 20 Optimization</b>	<b>158</b>
<b>CHAPTER 21 Risk Budgets</b>	<b>173</b>
<b>CHAPTER 22 Hedging</b>	<b>182</b>
<b>CHAPTER 23 Option Valuation and Replication</b>	<b>192</b>
<b>CHAPTER 24 Commodity Futures Contracts</b>	<b>201</b>
<b>CHAPTER 25 Currencies</b>	<b>210</b>
<b>Glossary</b>	<b>221</b>
<b>Notes</b>	<b>241</b>
<b>Index</b>	<b>253</b>

# Foreword

## Time and Magic: What This Book Is About

Once referred to financial markets as “dazzling creations.” Mark Kritzman’s superb essays in this book will show you where the dazzle is. Equally important, he will show you how to avoid being blinded by razzle-dazzle.

The book begins with a masterful summary of the contributions of the Nobel Prize winners to the theory of finance and investment. Before the 1950s, there was no theory in finance, just crude rules of thumb encrusted with fables, myths, and lore. Risk was measured in the gut; there were no rules for managing risk; expected return was what the latest tout proclaimed it to be; and the interactions among bonds, stocks, and other assets were a matter of little interest because bonds were for buy-and-hold investing and stocks were for trading; other assets were of no particular interest. In the course of about 20 years, from the early 1950s to the early 1970s, powerful advances in financial theory transformed a century or more of folklore about investing into a systematic structure for managing risk, for expanding the opportunities for higher rates of return, for calibrating performance, and for decision-making by issuers as well as buyers.

Financial theory as originally set forth by the Nobel laureates has undergone many modifications with the passage of time, not only in overcoming some of the oversimplifications that all theories exhibit, but also in broadening the applications of these ideas in ways no one ever conceived of in the early days. The art of investing will never be the same as it was before 1952.

This book demonstrates how investors can put these fundamental principles into practice. Nobody has made us so smart we can be certain of getting rich. But investors who heed the lessons here will gain a keen sense of how markets function; how to understand the essential character of different kinds of assets; how to measure, combine, and manage risks in a rational manner; and how to deal with the most puzzling and yet most dominant feature of investing: the passage of time.

Time is at the heart of every transaction in financial markets. Without financial markets, all assets would be buy-and-hold investments. Changing

your mind after you have made an investment, at a moment of your own choosing, would be extremely costly or even impossible under such conditions. You would hesitate a long time before buying an asset with these constraints, and even then you would demand high compensation for taking the risk that matters could take a turn for the worse without any opportunity for you to get out. Or perhaps more attractive investments might come along after you have committed your money to this illiquid situation. Let's face it: Many individuals in the investment management profession take a lot longer to reach a decision about which house to buy for their family, and how much to pay for it, than they spend before committing millions of dollars of client money in stocks like General Electric or Microsoft.

Financial markets are a kind of magic elixir providing you with the opportunity to change your mind after you have made a commitment. You may have to pay a price for that reversibility—a price you will never know in advance—but the trade-off for that risk is in sustaining your flexibility instead of locking yourself up in what might be a much greater risk over the indefinite future. In more technical terms, whatever the present value of the expected future cash flows from your investment might be, you can realize them in the markets at your discretion without having to wait through the future to cash out.

That is the picture from the perspective of an investor who wants to sell an asset, but someone else is always at the other side of every market transaction. If, as I have just pointed out, sellers are realizing the present value of future cash flows today, buyers are doing the opposite. They are using today's cash—or instant purchasing power—to purchase a stream of cash flows expected to arrive in the future. In short, *financial markets are a time machine that allows selling investors to compress the future into the present while allowing buying investors to stretch the present into the future.*

But there are no free lunches and no sure things. Through all of this wonderful magic lurks uncertainty. We do not know what the future holds, we never have known, and we never shall know. As Frank Sinatra used to phrase it, "That's life!" Everyone reading these words will nod their heads in agreement. But few people reading these words will *behave* as though the future is unknown. Consider what "the future is unknown" really means: that surprise is endemic, an inescapable feature of daily life. There would be no such thing as surprise if we knew the future.

Yet humans have never learned how to live with surprise, which is just another way of saying they persist in believing they know what lies ahead. The most vivid proof of that assertion is in stock markets, where volatility seems to be a permanent feature—sometimes more, sometimes less, but always present. These sharp and often discontinuous price movements in either direction are clear evidence that surprise takes us by surprise. If we

really accepted our ignorance of the future, the frequent arrival of the unexpected would not shake us up as much as it does. The very word “unexpected” shows we had a view of what the future would bring even though we admit we do not know what the future will bring. The eternal conflict between what we expect and what actually does happen is encapsulated in that important little word, risk.

The many wisdoms in this book focus primarily on that conflict. Mark Kritzman’s objective is to show you how you can deal with the harsh reality that we do not know what the outcomes of our decisions will be. Risk is an inescapable feature of investing, and risk, more than anything else, is what this fine book is all about.

PETER L. BERNSTEIN



# Preface

**T**he *Portable Financial Analyst* is a practitioner's guide to the important concepts, analytical methods, and investment strategies of modern finance. I have tried to present this material in a style accessible to most practitioners with minimal technical jargon and mathematical symbolism. Moreover, the content is reasonably self-contained, which means you need not refer to other sources in order to proceed comfortably through these pages. My goal is to facilitate understanding rather than to persuade you to switch careers.

What is different about this second edition? I have expanded some of the original chapters, added four new chapters, clarified and refined material as needed, added a glossary of technical terms, unified the mathematical notation, and corrected alleged errors in the first edition. Thankfully, most of the revisions are additions rather than corrections. On the other hand, corrections might motivate more people who own the first edition to purchase this new edition.

New chapters include:

- Higher Moments
- Event Studies
- Value at Risk
- Risk Budgets

Some of the major additions to the original material include:

- *Time Diversification*—a discussion of option valuation and first passage risk as challenges to the notion that time diversifies risk.
- *Simulation*—a description of bootstrapping simulation, along with an application to determine the hierarchy of investment choice.
- *Future Value and Risk of Loss*—a discussion of within-horizon exposure to loss, including its estimation and comparison to conventional risk measures.
- *Optimization*—a description of Chow's mean-variance-tracking error optimization and an exposé of the sophistry surrounding the tiresome critique, "garbage in, garbage out."

- *Option Valuation and Replication*—an expanded discussion of option valuation as preparation for the review of option replication.
- *Currencies*—an appendix detailing the mathematics of currency hedging.

I previewed some of the new material in Peter Bernstein's publication, *Economics and Portfolio Strategy*. As a consequence, I have benefited enormously from Peter's insights and editorial acumen. Should you discover anything approaching eloquence in these pages, more likely than not Peter is responsible.

I am grateful to Sebastien Page, who unified the mathematical notation. The notation in the first edition was quite disjointed, because the chapters originally appeared as separate articles in the *Financial Analysts Journal*.

I owe a special debt to George Chow, my business partner and friend, who has patiently endured many long discussions about these topics and who has originated some of the ideas presented herein.

I have benefited from conversations and correspondence with many acquaintances and colleagues, who have helped me to refine and advance my understanding of these topics. Although the complete list escapes my memory, several people stand out: Jeremy Armitage, Stephen Brown, Roger Clarke, Ken Froot, Gary Gastineau, Eric Lobben, Chip Lowry, Alan Marcus, Harry Markowitz, Jack Meyer, Paul O'Connell, Krishna Ramaswamy, Gita Rao, Don Rich, Paul A. Samuelson, Bill Sharpe, Stan Shelton, Jack Treynor, and Anne-Sophie Vanroyen.

I wish to acknowledge the continued support of the Association for Investment Management and Research and, in particular, Katy Sherrerd.

I am grateful to Pamela van Giessen, my editor at John Wiley & Sons, who has been a determined and patient advocate of this book. Her guidance, humor, and friendship have been a constant source of encouragement.

My final thanks are to my wife, Elizabeth Gorman, who graciously chose to pursue a Ph.D. while I worked on this book, thereby affording me much needed time for this task. Moreover, on her way to a Ph.D. she developed quite a proficiency in quantitative methods, which proved to be most beneficial. But, of course, I am most thankful for her encouragement, understanding, and sound counsel.

# The Nobel Prize

**O**n October 16, 1990, the Royal Swedish Academy of Sciences announced its selection for the Nobel Memorial Prize in Economic Science. For the first time since the prize for economics was established in 1968, the Royal Academy chose three individuals whose primary contributions were in finance and whose affiliations were not with arts and science schools, but rather with schools of business. Harry Markowitz was cited for his pioneering research in portfolio selection, while William Sharpe shared the award for developing an equilibrium theory of asset pricing. Merton Miller was a co-winner for his contributions in corporate finance, in which he showed, along with Franco Modigliani, that the value of a firm should be invariant to its capital structure and dividend policy.

The pioneering research of these individuals revolutionized finance and accelerated the application of quantitative methods to financial analysis.

## **PORTFOLIO SELECTION**

---

In his classic article, “Portfolio Selection,” Markowitz submitted that investors should not choose portfolios that maximize expected return, because this criterion by itself ignores the principle of diversification.<sup>1</sup> He proposed that investors should instead consider variances of return, along with expected returns, and choose portfolios that offer the highest expected return for a given level of variance. He called this rule the E-V maxim.

Markowitz showed that a portfolio’s expected return is simply the weighted average of the expected returns of its component securities. A portfolio’s variance is a more complicated concept, however. It depends on more than just the variances of the component securities.

The variance of an individual security is a measure of the dispersion of its returns. It is calculated by squaring the difference between each return in a series and the mean return for the series, then averaging these squared differences. The square root of the variance (the standard deviation) is often

used in practice because it measures dispersion in the same units in which the underlying return is measured.

Variance provides a reasonable gauge of a security's risk, but the average of the variances of two securities will not necessarily give a good indication of the risk of a portfolio comprising these two securities. The portfolio's risk depends also on the extent to which the two securities move together—that is, the extent to which their prices react in like fashion to a particular event.

To quantify co-movement among security returns, Markowitz introduced the statistical concept of covariance. The covariance between two securities equals the standard deviation of the first times the standard deviation of the second times the correlation between the two.

The correlation, in this context, measures the association between the returns of two securities. It ranges in value from 1 to  $-1$ . If one security's returns are higher than its average return when another security's returns are higher than its average return, for example, the correlation coefficient will be positive, somewhere between 0 and 1. Alternatively, if one security's returns are lower than its average return when another security's returns are higher than its average return, then the correlation will be negative.

The correlation, by itself, is an inadequate measure of covariance because it measures only the direction and degree of association between securities' returns. It does not account for the magnitude of variability in each security's returns. Covariance captures magnitude by multiplying the correlation by the standard deviations of the securities' returns.

Consider, for example, the covariance of a security with itself. Obviously, the correlation in this case equals 1. A security's covariance with itself thus equals the standard deviation of its returns squared, which, of course, is its variance.

Finally, portfolio variance depends also on the weightings of its constituent securities—the proportion of a portfolio's market value invested in each. The variance of a portfolio consisting of two securities equals the variance of the first security times its weighting squared plus the variance of the second security times its weighting squared plus twice the covariance between the securities times each security's weighting. The standard deviation of this portfolio equals the square root of the variance.

From this formulation of portfolio risk, Markowitz was able to offer two key insights. First, unless the securities in a portfolio are perfectly inversely correlated (that is, have a correlation of  $-1$ ), it is not possible to eliminate portfolio risk entirely through diversification. If we divide a portfolio equally among its component securities, for example, as the number of securities in the portfolio increases, the portfolio's risk will

tend not toward zero but, rather, toward the average covariance of the component securities.

Second, unless all the securities in a portfolio are perfectly positively correlated with each other (a correlation of 1), a portfolio's standard deviation will always be less than the weighted average standard deviation of its component securities. Consider, for example, a portfolio consisting of two securities, both of which have expected returns of 10 percent and standard deviations of 20 percent and which are uncorrelated with each other. If we allocate portfolio assets equally between these two securities, the portfolio's expected return will equal 10 percent, while its standard deviation will equal 14.14 percent. The portfolio offers a reduction in risk of nearly 30 percent relative to investment in either of the two securities separately. Moreover, this risk reduction is achieved without any sacrifice of expected return.

Markowitz also demonstrated that, for given levels of risk, we can identify particular combinations of securities that maximize expected return. He deemed these portfolios "efficient" and referred to a continuum of such portfolios in dimensions of expected return and standard deviation as the efficient frontier. According to Markowitz's E-V maxim, investors should choose portfolios located along the efficient frontier. It is almost always the case that there exists some portfolio on the efficient frontier that offers a higher expected return and less risk than the least risky of its component securities (assuming the least risky security is not completely riskless).

The financial community was slow to implement Markowitz's theory, in large part because of a practical constraint. In order to estimate the risk of a portfolio of securities, one must estimate the variances of every security, along with the covariances between every pair of securities. For a portfolio of 100 securities, this means calculating 100 variances and 4,950 covariances—5,050 risk estimates! In general, the number of required risk estimates (variances and covariances) equals  $n(n + 1)/2$ , where  $n$  equals the number of securities in the portfolio.<sup>2</sup> In 1952, when Markowitz published "Portfolio Selection," the sheer number of calculations formed an obstacle in the way of acceptance. It was in part the challenge of this obstacle that motivated William Sharpe to develop a single index measure of a security's risk.

## **THE CAPITAL ASSET PRICING MODEL**


James Tobin, the 1981 winner of the Nobel Prize in economics, showed that the investment process can be separated into two distinct steps: (1) the construction of an efficient portfolio, as described by Markowitz, and (2)

the decision to combine this efficient portfolio with a riskless investment. This two-step process is the famed separation theorem.<sup>3</sup>

Sharpe extended Markowitz's and Tobin's insights to develop a theory of market equilibrium under conditions of risk.<sup>4</sup> First, Sharpe showed that there is along the efficient frontier a unique portfolio that, when combined with lending or borrowing at the riskless interest rate, dominates all other combinations of efficient portfolios and lending or borrowing.

Figure 1.1 shows a two-dimensional graph, with risk represented by the horizontal axis and expected return represented by the vertical axis. The efficient frontier appears as the positively sloped concave curve. The straight line emanating from the vertical axis at the riskless rate illustrates the efficient frontier with borrowing and lending. The segment of the line between the vertical axis and the efficient portfolio curve represents some combination of the efficient portfolio M and lending at the riskless rate, while points along the straight line to the right represent some combination of the efficient portfolio and borrowing at the riskless rate. Combinations of portfolio M and lending or borrowing at the riskless rate will always offer the highest expected rate of return for a given level of risk.

With two assumptions, Sharpe demonstrated that in equilibrium investors will prefer points along the line emanating from the riskless rate that is tangent to M. The requisite assumptions are (1) there exists a single



**FIGURE 1.1** Efficient frontier with borrowing and lending

riskless rate at which investors can lend and borrow in unlimited amounts, and (2) investors have homogeneous expectations regarding expected returns, variances, and covariances. Under these assumptions, Sharpe showed that portfolio M is the market portfolio, which represents the maximum achievable diversification.

Within this model, Sharpe proceeded to demonstrate that risk can be partitioned into two sources—that caused by changes in the value of the market portfolio, which cannot be diversified away, and that caused by nonmarket factors, which is diversified away in the market portfolio. He labeled the nondiversifiable risk systematic risk and the diversifiable risk unsystematic risk.

Sharpe also showed that a security's systematic risk can be estimated by regressing its returns (less the riskless rate) against the market portfolio's returns (less riskless rate). The slope from this regression equation, which Sharpe called beta, quantifies the security's systematic risk when multiplied by the market risk. The unexplained variation in the security's return (the residuals from the regression equation) represents the security's unsystematic risk. He then asserted that, in an efficient market, investors are only compensated for bearing systematic risk, because it cannot be diversified away, and the expected return of a security is, through beta, linearly related to the market's expected return.

It is important to distinguish between a single index model and the Capital Asset Pricing Model (CAPM). A single index model does not require the intercept of the regression equation (alpha) to equal 0 percent. It simply posits a single source of systematic, or common, risk. Stated differently, the residuals from the regression equation are uncorrelated with each other. The important practical implication is that it is not necessary to estimate covariances between securities. Each security's contribution to portfolio risk is captured through its beta coefficient. The CAPM, by contrast, does require the intercept of the regression equation to equal 0 percent in an efficient market. The CAPM itself does not necessarily assume a single source of systematic risk. This is tantamount to allowing for some correlation among the residuals.

## **INVARIANCE PROPOSITIONS**

---

Between the publication of Markowitz's theory of portfolio selection and Sharpe's equilibrium theory of asset pricing, Franco Modigliani (the 1985 Nobel Prize winner in economics) and Merton Miller published two related articles in which they expounded their now famous invariance propositions. The first, "The Cost of Capital, Corporation Finance, and the Theory of Investment," appeared in 1958.<sup>5</sup> It challenged the then

conventional wisdom that a firm's value depends on its capital structure (i.e., its debt/equity mix).

In challenging this traditional view, Modigliani and Miller invoked the notion of arbitrage. They argued that if a leveraged firm is undervalued, investors can purchase its debt and its shares. The interest paid by the firm is offset by the interest received by the investors, so the investors end up holding a pure equity stream. Alternatively, if an unleveraged firm is undervalued, investors can borrow funds to purchase its shares. The substitutability of individual debt for corporate debt guarantees that firms in the same risk class will be valued the same, regardless of their respective capital structures. In essence, Modigliani and Miller argued in favor of the law of one price.

In a subsequent article, "Dividend Policy, Growth, and the Valuation of Shares," Modigliani and Miller proposed that a firm's value is invariant, not only to its capital structure, but also to its dividend policy (assuming the firm's investment decision is set independently).<sup>6</sup> Again, they invoked the notion of substitutability, arguing that repurchasing shares has the same effect as paying dividends; thus issuing shares and paying dividends is a wash. Although the cash component of an investor's return may differ as a function of dividend policy, the investor's total return, including price change, should not change with dividend policy.

Modigliani and Miller's invariance propositions provoked an enormous amount of debate and research. Much of the sometimes spirited debate centered on the assumption of perfect capital markets. In the real world, where investors cannot borrow and lend at the riskless rate of interest, where both corporations and individuals pay taxes, and where investors do not share equal access with management to relevant information, there is only spotty evidence to support Modigliani and Miller's invariance propositions.

But the value of the contributions of these Nobel laureates does not depend on the degree to which their theories hold in an imperfect market environment. It depends, rather, on the degree to which they changed the financial community's understanding of the capital markets. Markowitz taught us how to evaluate investment opportunities probabilistically, while Sharpe provided us with an equilibrium theory of asset pricing, enabling us to distinguish between risk that is rewarded and risk that is not rewarded. Miller, in collaboration with Modigliani, demonstrated how the simple notion of arbitrage can be applied to determine value, which subsequently was extended to option valuation—yet another innovation that proved worthy of the Nobel Prize.

# Uncertainty

**T**he primary challenge to financial analysts is to determine how to proceed in the face of uncertainty. Uncertainty arises from imperfect knowledge and from incomplete data. Methods for interpreting limited information may thus help analysts measure and control uncertainty.

Long ago, natural scientists noticed the widespread presence of random variation in nature. This led to the development of laws of probability, which help predict outcomes. As it turns out, many of the laws that seem to explain the behavior of random variables in nature apply as well to the behavior of financial variables such as corporate earnings, interest rates, and asset prices and returns.

## **RELATIVE FREQUENCY**

---

A random variable can be thought of as an event whose outcomes in a given situation depend on chance factors. For example, the toss of a coin is an event whose outcome is governed by chance, as is next year's closing price for the stock market. Because an outcome is influenced by chance does not mean we are completely ignorant about its possible values. We may, for example, be able to garner some insights from prior experiences.

Suppose we are interested in predicting the return of the stock market over the next 12 months. Should we be more confident in predicting that it will be between 0 and 10 percent than between 10 and 20 percent? The past history of stock market returns can tell us how often returns within specified ranges have occurred. Table 2.1 shows annual stock market returns over a 40-year period.

We simply count the number of returns between 0 and 10 percent and the number of returns between 10 and 20 percent. Dividing each figure by 40 gives us the relative frequency of returns within each range. Six returns fall within the range of 0 to 10 percent, while 10 returns fall within the

**TABLE 2.1** Annual Stock Market Returns

Year	Return	Year	Return	Year	Return	Year	Return
1	24.00%	11	26.90%	21	14.30%	31	-4.90%
2	18.40%	12	-8.70%	22	19.00%	32	21.40%
3	-1.00%	13	22.80%	23	-14.70%	33	22.50%
4	52.60%	14	16.50%	24	-26.50%	34	6.30%
5	31.60%	15	12.50%	25	37.20%	35	32.20%
6	6.60%	16	-10.10%	26	23.80%	36	18.80%
7	-10.80%	17	24.00%	27	-7.20%	37	5.30%
8	43.40%	18	11.10%	28	6.60%	38	16.60%
9	12.00%	19	-8.50%	29	18.40%	39	31.80%
10	0.50%	20	4.00%	30	32.40%	40	-3.10%

range of 10 to 20 percent. The relative frequencies of these observations are 15 and 25 percent, respectively, as Table 2.2 shows.

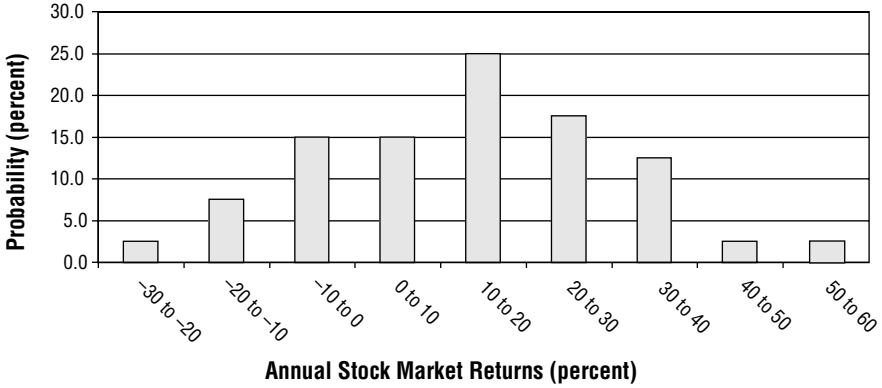
Figure 2.1 depicts this information graphically in what is called a discrete probability distribution. (It is discrete because it covers a finite number of observations rather than an infinite number of continuously distributed observations.) The values along the vertical axis represent the probability (equal here to the relative frequency) of observing a return within the ranges indicated along the horizontal axis.

The information we have is limited. For one thing, the return ranges (which we set) are fairly wide. For another, the sample is confined to annual returns and covers only a particular 40-year period, which may exclude important events such as wars or perhaps the Great Depression.

We can nonetheless draw some tentative inferences from this limited information. For example, we may assume we are about two-thirds more

**TABLE 2.2** Frequency Distribution

Range of Return	Frequency	Relative Frequency
-30% to -20%	1	2.5%
-20% to -10%	3	7.5%
-10% to 0%	6	15.0%
0% to 10%	6	15.0%
10% to 20%	10	25.0%
20% to 30%	7	17.5%
30% to 40%	5	12.5%
40% to 50%	1	2.5%
50% to 60%	1	2.5%



**FIGURE 2.1** Discrete probability distribution

likely to observe a return within the range of 10 to 20 percent than a return within the range of 0 to 10 percent. Furthermore, by summing the relative frequencies for the three ranges below 0 percent, we may also assume there is a 25 percent chance of experiencing a negative return.

If we want to draw more precise inferences, we should increase the sample size and partition the data into narrower ranges. If we proceed along these lines, the distribution of returns should eventually resemble the familiar pattern known as the bell-shaped curve, or normal distribution.

## **NORMAL DISTRIBUTION**

The normal distribution is a continuous probability distribution; it assumes there are an infinite number of observations covering all possible values along a continuous scale. Time, for example, can be thought of as being distributed along a continuous scale. Stocks, however, trade in discrete units, so technically stock returns cannot be distributed continuously. Nonetheless, for purposes of financial analysis, the normal distribution is usually a reasonable approximation of the distribution of stock returns, as well as the returns of other financial assets.

The formula that gives rise to the normal distribution was first published by Abraham de Moivre in 1733. Its properties were investigated by Carl Gauss in the 18th and 19th centuries. In recognition of Gauss's contributions, the normal distribution is often referred to as the Gaussian distribution.

The normal distribution has special appeal to natural scientists for two reasons. First, it is an excellent approximation of the random variation of

many natural phenomena.<sup>1</sup> Second, it can be described fully by only two values: (1) the mean of the observations, which measures location or central tendency, and (2) the variance of the observations, which measures dispersion.

For our sample of annual stock market returns, the mean return (which is also the expected return) equals the sum of the observed returns times their probabilities of occurrence:

$$\mu = \text{Pr}_1 \times R_1 + \text{Pr}_2 \times R_2 + \dots + \text{Pr}_n \times R_n \quad (2.1)$$

where  $\mu$  = the mean return  
 $R_1, R_2 \dots R_n$  = observed returns in years 1 through  $n$   
 $\text{Pr}_1, \text{Pr}_2 \dots \text{Pr}_n$  = the probabilities of occurrence (or relative frequencies) of the returns in years 1 through  $n$

The variance of returns ( $\sigma^2$ ) is computed as the probability-weighted squared difference from the mean. To compute the variance, we subtract the mean return from each annual return, square these values, sum these squared values, and then weight them by their probability of occurrence.

$$\sigma^2 = \text{Pr}_1(R_1 - \mu) + \text{Pr}_2(R_2 - \mu) + \dots + \text{Pr}_n(R_n - \mu) \quad (2.2)$$

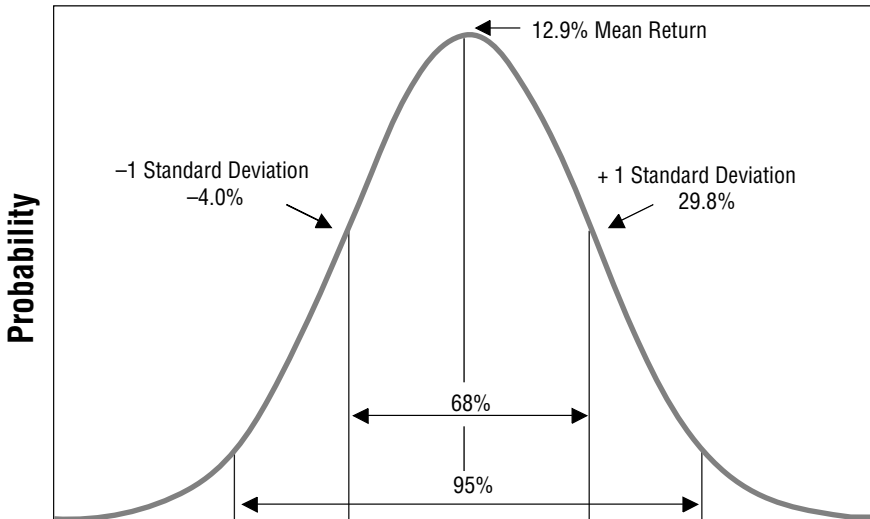
$\text{Pr}_1, \text{Pr}_2 \dots \text{Pr}_n$  = the probabilities of occurrence (or relative frequencies) of the returns in years 1 through  $n$

The square root of the variance, which is called the standard deviation, is commonly used as a measure of dispersion.

If we apply these formulas to the annual returns in Table 2.1, we find that the mean return for the sample equals 12.9 percent, the variance of returns equals 2.9 percent, and the standard deviation of returns equals 16.9 percent. These values, together with the assumption that the returns of the stock market are normally distributed, enable us to infer a normal probability distribution of stock market returns. This probability distribution is shown in Figure 2.2.

The normal distribution has several important characteristics. First, it is symmetric around its mean; 50 percent of the returns are below the mean return, and 50 percent of the returns are above the mean return. Also, because of this symmetry, the mode of the sample (the most common observation) and the median (the middle value of the observations) are equal to each other and to the mean.

Note that the area enclosed within one standard deviation on either side of the mean encompasses 68 percent of the total area under the curve. The area enclosed within two standard deviations on either side of the



**FIGURE 2.2** Normal probability distribution

mean encompasses 95 percent of the total area under the curve, and more than 99 percent of the area under the curve falls within plus and minus three standard deviations of the mean.

From this information we are able to draw several conclusions. For example, we know that 68, 95, and more than 99 percent of returns, respectively, will fall within one, two, and three standard deviations (plus and minus) of the mean return. It is thus straightforward to measure the probability of experiencing returns that are one, two, or three standard deviations away from the mean.

There is, for example, about a 32 percent chance (100 percent minus 68 percent) of experiencing returns at least one standard deviation above or below the mean return. Thus there is only a 16 percent chance of experiencing a return below  $-4.0$  percent (mean of  $12.9$  percent minus standard deviation of  $16.9$  percent) and an equal chance of experiencing a return greater than  $29.8$  percent (mean of  $12.9$  percent plus standard deviation of  $16.9$  percent).

## **STANDARDIZED VARIABLES**

---

We may, however, be interested in the likelihood of experiencing a return less than  $0$  percent or a return greater than  $15$  percent. In order to determine the

probabilities of these returns (or the probability of achieving any return, for that matter), we can standardize the target return. We do so by subtracting the mean return from the target return and dividing by the standard deviation. (By standardizing returns we, in effect, rescale the distribution to have a mean of 0 and a standard deviation of 1.) Thus, to find the area under the curve to the left of 0 percent (which is the same as the probability of experiencing a return of less than 0 percent), we subtract 12.9 percent (the mean) from 0 percent (the target) and divide this quantity by 16.9 percent (the standard deviation):

$$\frac{0.0 - 0.129}{0.169} = -0.7633$$

This value tells us that 0 percent is 0.7633 standard deviations below the mean. This is much less than a full standard deviation, so we know that the chance of experiencing a return of less than 0 percent must be greater than 16 percent.

In order to calculate a precise probability directly, we need to evaluate the integral of the standardized normal density function. Fortunately, most statistics books and spreadsheet software show the area under a standardized normal distribution curve that corresponds to a particular standardized variable. Table 2.3 is one example.

To find the area under the curve to the left of the standardized variable, we read down the left column to the value  $-0.7$  and across this row to the column under the value  $-0.06$ . The value at this location, 0.2236, equals the probability of experiencing a return of less than 0 percent. This, of course, implies that the chance of experiencing a return greater than 0 percent equals 0.7764 (i.e.,  $1 - 0.2236$ ). By comparison, the probability of experiencing a negative return as estimated from the discrete probability distribution in Table 2.2 equals 25 percent.

Suppose we are interested in the likelihood of experiencing an annualized return of less than 0 percent on average over a five-year horizon? First, we'll assume the year-to-year returns are mutually independent (that is, this year's return has no effect on next year's return). We convert the standard deviation back to the variance (by squaring it), divide the variance by 5 (the number of years in the horizon), and use the square root of this value to standardize the difference between 0 percent and the mean return. Alternatively, we can simply divide the standard deviation by the square root of 5 and use this value to standardize the difference:

$$\frac{0.0 - 0.129}{0.169 \times \sqrt{5}} = -1.71$$

Again, by referring to Table 2.3, we find that the likelihood of experiencing an annualized return less than 0 percent on average over five years equals only 0.0436, or 4.36 percent. This is much less than the probability of experiencing a negative return in any one year. Intuitively, we are less likely to lose money on average over five years than in any particular year because we are diversifying across time; a loss in any particular year might be offset by a gain in one or more of the other years.

Now suppose we are interested in the likelihood that we might lose money in one or more of the five years. This probability is equivalent to one minus the probability of experiencing a positive return in every one of the five years. Again, if we assume independence in the year-to-year returns, the likelihood of experiencing five consecutive yearly returns each greater than 0 percent equals 0.7764 raised to the fifth power, which is 0.2821. Thus the probability of experiencing a negative return in at least one of the five years equals 0.7179 (i.e.,  $1 - 0.2821$ ).

Over extended holding periods, the normal distribution may not be a good approximation of the distribution of returns because short holding-period returns are compounded, rather than cumulated, to derive long holding-period returns. Because we can represent the compound value of an index as a simple accumulation when expressed in terms of logarithms, it is the logarithms of 1 plus the holding-period returns that are normally distributed. The actual returns thus conform to a lognormal distribution. A lognormal distribution assigns higher probabilities to extremely high values than it does to extremely low values; the result is a skewed distribution, rather than a symmetric one. This distinction is usually not significant for holding periods of one year or less. For longer holding periods, the distinction can be important. For this reason, we should assume a lognormal distribution when estimating the probabilities associated with outcomes over long investment horizons.<sup>2</sup>

## **CAVEATS**

---

In applying the normal probability distribution to measure uncertainty in financial analysis, we should proceed with caution. We must recognize, for example, that our probability estimates are subject to sampling error. Our example assumed implicitly that the 40-year sample characterized the mean and variance of stock market returns. This sample, in fact, represents but a small fraction of the entire universe of historical returns and may not necessarily be indicative of the central tendency and dispersion of returns going forward.

As an alternative to extrapolating historical data, we can choose to estimate the expected stock market return based on judgmental factors, and

**TABLE 2.3** Normal Distribution Table

Probability That Standardized Variable Is Less Than Z												
Z	0	-0.01	-0.02	-0.03	-0.04	-0.05	-0.06	-0.07	-0.08	-0.09		
0.0	0.5000	0.4960	0.4920	0.4880	0.4840	0.4785	0.4761	0.4721	0.4681	0.4641		
-0.1	0.4602	0.4562	0.4522	0.4483	0.4443	0.4388	0.4364	0.4325	0.4286	0.4247		
-0.2	0.4207	0.4168	0.4129	0.4090	0.4052	0.3997	0.3974	0.3936	0.3897	0.3859		
-0.3	0.3821	0.3783	0.3745	0.3707	0.3669	0.3617	0.3594	0.3557	0.3520	0.3483		
-0.4	0.3446	0.3409	0.3372	0.3336	0.3300	0.3249	0.3228	0.3192	0.3156	0.3121		
-0.5	0.3085	0.3050	0.3015	0.2981	0.2946	0.2898	0.2877	0.2843	0.2810	0.2776		
-0.6	0.2743	0.2709	0.2676	0.2643	0.2611	0.2566	0.2546	0.2514	0.2483	0.2451		
-0.7	0.2420	0.2389	0.2358	0.2327	0.2296	0.2254	0.2236	0.2206	0.2177	0.2148		
-0.8	0.2119	0.2090	0.2061	0.2033	0.2005	0.1966	0.1949	0.1922	0.1894	0.1867		
-0.9	0.1841	0.1814	0.1788	0.1762	0.1736	0.1700	0.1685	0.1660	0.1635	0.1611		
-1.0	0.1587	0.1562	0.1539	0.1515	0.1492	0.1459	0.1446	0.1423	0.1401	0.1379		
-1.1	0.1357	0.1335	0.1314	0.1292	0.1271	0.1243	0.1230	0.1210	0.1190	0.1170		
-1.2	0.1151	0.1131	0.1112	0.1093	0.1075	0.1049	0.1038	0.1020	0.1003	0.0985		
-1.3	0.0968	0.0951	0.0934	0.0918	0.0901	0.0879	0.0869	0.0853	0.0838	0.0823		