

Wiley Finance Series

Quantitative Financial Risk Management

MICHAEL B. MILLER

WILEY

QUANTITATIVE FINANCIAL RISK MANAGEMENT

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.

QUANTITATIVE FINANCIAL RISK MANAGEMENT

Michael B. Miller

WILEY

Copyright © 2019 by Michael B. Miller. 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) 646-8600, 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 www.wiley.com/go/permissions.

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 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 publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Cataloging-in-Publication Data

Names: Miller, Michael B. (Michael Bernard), 1973- author.

Title: Quantitative financial risk management / Michael B. Miller.

Description: Hoboken, New Jersey : Wiley, [2019] | Series: Wiley finance

series | Includes bibliographical references and index. |

Identifiers: LCCN 2018033207 (print) | LCCN 2018044462 (ebook) | ISBN

9781119522232 (Adobe PDF) | ISBN 9781119522263 (ePub) | ISBN 9781119522201

| ISBN 9781119522201 (hardcover) | ISBN 9781119522232 (ePDF) | ISBN

9781119522263 (ePub)

Subjects: LCSH: Financial risk management.

Classification: LCC HD61 (ebook) | LCC HD61 .M5373 2019 (print) | DDC

332—dc23

LC record available at <https://lcn.loc.gov/2018033207>

Cover Design: Wiley

Cover Images: © Sergey Nivens/Shutterstock; © whiteMocca/Shutterstock

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

Preface	vii
About the Author	ix
1 Overview of Financial Risk Management	1
2 Market Risk: Standard Deviation	15
3 Market Risk: Value at Risk	51
4 Market Risk: Expected Shortfall, and Extreme Value Theory	73
5 Market Risk: Portfolios and Correlation	91
6 Market Risk: Beyond Correlation	119
7 Market Risk: Risk Attribution	151
8 Credit Risk	167
9 Liquidity Risk	189
10 Bayesian Analysis	205
11 Behavioral Economics and Risk	231
Appendix A Maximum Likelihood Estimation	247
Appendix B Copulas	253
Answers to End-of-Chapter Questions	257
References	295
Index	297

PREFACE

My first book on financial risk management, *Mathematics and Statistics for Financial Risk Management*, grew out of my experience working in the hedge fund industry and my involvement with the Global Association of Risk Professionals. It was written for practitioners who may not have had the opportunity to take the advanced courses in mathematics— especially those courses in statistics—that are necessary for a deeper understanding of modern financial risk management. It was also for practitioners who had taken these courses but may have forgotten what they learned. To be honest, I often use the first book as a reference myself. Even authors forget.

As a result of that first book, I was asked to teach a graduate-level course in risk management. I realized that my students had the opposite problem of my colleagues in the hedge fund industry. My students came to the course with a very strong foundation in mathematics, but knew less about the workings of financial markets or the role of risk managers within a financial firm. This book was written for them, and I have been teaching with the material that this book is based on for a number of years now.

There is considerable overlap between the two books. Indeed, there are some sections that are almost identical. While the first book was organized around topics in mathematics, however, this book is organized around topics in risk management. In each chapter we explore a particular topic in risk management along with various mathematical tools that can be used to understand that topic. As with the first book, I have tried to provide a large number of sample problems and practical end-of-chapter questions. I firmly believe that the best way to understand financial models is to work through actual problems.

This book assumes that the reader is familiar with basic calculus, linear algebra, and statistics. When a particular topic in mathematics is central to a topic in risk management, I review the basics and introduce notation, but the pace can be quick. For example, in the first chapter we review standard deviation, but we only spend one section on what would likely be an entire chapter in an introductory book on statistics.

Risk management in practice often requires building models using spreadsheets or other financial software. Many of the topics in this book are accompanied by an icon, shown here:



These icons indicate that Excel examples can be found at John Wiley & Sons' companion website for *Quantitative Financial Risk Management*, www.wiley.com/go/millerfinancialrisk.

ABOUT THE AUTHOR

Michael B. Miller is the founder and CEO of Northstar Risk Corp. Before starting Northstar, Mr. Miller was Chief Risk Officer for Tremblant Capital and, before that, Head of Quantitative Risk Management at Fortress Investment Group.

Mr. Miller is the author of *Mathematics and Statistics for Financial Risk Management*, now in its second edition, and, along with Emanuel Derman, *The Volatility Smile*. He is also an adjunct professor at Columbia University and the co-chair of the Global Association of Risk Professional's Research Fellowship Committee. Before starting his career in finance, Mr. Miller studied economics at the American University of Paris and the University of Oxford.

QUANTITATIVE FINANCIAL RISK MANAGEMENT

1

OVERVIEW OF FINANCIAL RISK MANAGEMENT

Imagine you are a chef at a restaurant. You've just finished preparing eggs benedict for a customer. The eggs are cooked perfectly, the hollandaise sauce has just the right mix of ingredients, and it all sits perfectly on the plate. The presentation is perfect! You're so proud of the way this has turned out that you decide to deliver the dish to the customer yourself. You place the plate in front of the customer, and she replies, "This looks great, but I ordered a filet mignon, and you forgot my drink."

Arguably, the greatest strength of modern financial risk management is that it is highly objective. It takes a scientific approach, using math and statistics to measure and evaluate financial products and portfolios. While these mathematical tools can be very powerful, they are simply that—tools. If we make unwarranted assumptions, apply models incorrectly, or present results poorly—or if our findings are ignored—then the most elegant mathematical models in the world will not help us. The eggs might be perfect, but that's irrelevant if the customer ordered a steak.

This is not a new idea, Vitruvius, a famous Roman architect wrote, "*Neque enim ingenium sine disciplina aut disciplina sine ingenio perfectum artificem potest efficere*", which roughly translates to "Neither genius without knowledge, nor knowledge without genius, will make a perfect artist." Applying this to risk management, we might say, "Neither math without knowledge of financial markets, nor knowledge of financial markets without math, will make a perfect risk manager."

Before we get to the math and statistics, then, we should take a step back and look at risk management more broadly. Before delving into the models, we explore the following

questions: What is risk management? What is the proper role for a risk manager within a financial organization? What do risk managers actually do on a day-to-day basis?

We end this chapter with a brief history of risk management. As you will see, risk management has made many positive contributions to finance, but it is far from being a solved problem.

WHAT IS RISK?

Before we can begin to describe what financial risk managers *do*, we need to understand what financial risk *is*. In finance, risk arises from uncertainty surrounding future profits or returns. There are many ways to define risk, and we may change the definition slightly, depending on the task at hand.

In everyday speech, the word *risk* is associated with the possibility of negative outcomes. For something to be risky, the final outcome must be uncertain *and* there must be some possibility that the final outcome will have negative consequences. While this may seem obvious, some popular risk measures treat positive and negative outcomes equally, while others focus only on negative outcomes. For this reason, in order to avoid any ambiguity when dealing specifically with negative outcomes, risk managers will often talk about *downside risk*.

Risk is often defined relative to expectations. If we have one investment with a 50/50 chance of earning \$0 or \$200, and a second investment with a 50/50 chance of earning \$400 or \$600, are both equally risky? The first investment earns \$100 on average, and the second \$500, but both have a 50/50 chance of being \$100 above or below this expected value. Because the deviations from expectations are equal, many risk managers would consider the two investments to be equally risky. By this logic, the second investment is more attractive because it has a higher expected return, not because it is less risky.

It is also important to note that risk is about *possible* deviations from expectations. If we expect an investment to make \$1 and it does make \$1, the investment was not necessarily risk free. If there were any possibility that the outcome could have been something other than \$1, then the investment was risky.

Absolute, Relative, and Conditional Risk

There may be no better way to understand the limits of financial risk management—why and where it may fail or succeed—than to understand the difference between absolute, relative, and conditional risk.

Financial risk managers are often asked to assign probabilities to various financial outcomes. What is the probability that a bond will default? What is the probability that an

equity index will decline by more than 10% over the course of a year? These types of predictions, where risk managers are asked to assess the total or *absolute risk* of an investment, are incredibly difficult to make. As we will see, assessing the accuracy of these types of predictions, even over the course of many years, can be extremely difficult.

It is often much easier to determine relative risk than to measure risk in isolation. Bond ratings are a good example. Bond ratings can be used to assess absolute risk, but they are on much surer footing when used to assess *relative risk*. The number of defaults in a bond portfolio might be much higher or lower next year depending on the state of the economy and financial markets. No matter what happens, though, a portfolio consisting of a large number of AAA-rated bonds will almost certainly have fewer defaults than a portfolio consisting of a large number of C-rated bonds. Similarly, it is much easier to say that emerging market equities are riskier than U.S. equities, or that one hedge fund is riskier than another hedge fund.

What is the probability that the S&P 500 will be down more than 10% next year? What is the probability that a particular U.S. large-cap equity mutual fund will be down more than 8% next year? Both are very difficult questions. What is the probability that this same mutual fund will be down more than 8%, *if* the S&P 500 is down more than 10%? This last question is actually much easier to answer. What's more, these types of *conditional risk* forecasts immediately suggest ways to hedge and otherwise mitigate risk.

Given the difficulty of measuring absolute risk, risk managers are likely to be more successful if they limit themselves to relative and conditional forecasts, when possible. Likewise, when there is any ambiguity about how a risk measure can be interpreted—as with bond ratings—encouraging a relative or conditional interpretation is likely to be in a risk manager's best interest.

Intrinsic and Extrinsic Risk

Some financial professionals talk about *risk* versus *uncertainty*. A better approach might be to contrast *intrinsic* risk and *extrinsic* risk.

When evaluating financial instruments, there are some risks that we consider to be intrinsic. No matter how much we know about the financial instrument we are evaluating, there is nothing we can do to reduce this intrinsic risk (other than reducing the size of our investment).

In other circumstances risk is due only to our own ignorance. In theory, this extrinsic risk can be eliminated by gathering additional information through research and analysis.

As an example, an investor in a hedge fund may be subject to both extrinsic and intrinsic risk. A hedge fund investor will typically not know the exact holdings of a hedge fund in which they are invested. Not knowing what securities are in a fund is extrinsic risk.

For various reasons, the hedge fund manager may not want to reveal the fund's holdings, but, at least in theory, this extrinsic risk *could be* eliminated by revealing the fund's holdings to the investor. At the same time, even if the investor *did* know what securities were in the fund, the returns of the fund would still not be fully predictable because the returns of the securities in the fund's portfolio are inherently uncertain. This inherent uncertainty of the security returns represents intrinsic risk and it cannot be eliminated, no matter how much information is provided to the investor.

Interestingly, a risk manager could reduce a hedge fund investor's extrinsic risk by explaining the hedge fund's risk guidelines. The risk guidelines could help the investor gain a better understanding of what might be in the fund's portfolio, without revealing the portfolio's precise composition.

Differentiating between these two fundamental types of risk is important in financial risk management. In practice, financial risk management is as much about reducing extrinsic risk as it is about managing intrinsic risk.

Risk and Standard Deviation

At the start of this chapter, we said that risk could be defined in terms of possible deviations from expectations. This definition is very close to the definition of standard deviation in statistics. The variance of a random variable is the expected value of squared deviations from the mean, and standard deviation is just the square root of variance. This is indeed very close to our definition of risk, and in finance risk is often equated with standard deviation.

While the two definitions are similar, they are not exactly the same. Standard deviation only describes what we expect the deviations will look like on average. Two random variables can have the same standard deviation, but very different return profiles. As we will see, risk managers need to consider other aspects of the distribution of expected deviations, not just its standard deviation.

WHAT IS FINANCIAL RISK MANAGEMENT?

In finance and in this book, we often talk about *risk management*, when it is understood that we are talking about *financial risk management*. Risk managers are found in a number of fields outside of finance, including engineering, manufacturing, and medicine.

When civil engineers are designing a levee to hold back flood waters, their risk analysis will likely include a forecast of the distribution of peak water levels. An engineer will often describe the probability that water levels will exceed the height of the levee in terms similar to those used by financial risk managers to describe the probability that losses in a portfolio

will exceed a certain threshold. In manufacturing, engineers will use risk management to assess the frequency of manufacturing defects. Motorola popularized the term Six Sigma to describe its goal to establish a manufacturing process where manufacturing defects were kept below 3.4 defects per million. (Confusingly the goal corresponds to 4.5 standard deviations for a normal distribution, not 6 standard deviations, but that's another story.) Similarly, financial risk managers will talk about big market moves as being three-sigma events or six-sigma events. Other areas of risk management can be valuable sources of techniques and terminology for financial risk management.

Within this broader field of risk management, though, how do we determine what is and is not financial risk management? One approach would be to define risk in terms of organizations, to say that *financial* risk management concerns itself with the risk of *financial* firms. By this definition, assessing the risks faced by Goldman Sachs or a hedge fund is financial risk management, whereas assessing the risks managed by the Army Corps of Engineers or NASA is not. A clear advantage to this approach is that it saves us from having to create a long list of activities that are the proper focus of financial risk management. The assignment is unambiguous. If a task is being performed by a financial firm, it is within the scope of financial risk management. This definition is future proof as well. If HSBC, one of the world's largest financial institutions, starts a new business line tomorrow, we do not have to ask ourselves if this new business line falls under the purview of financial risk management. Because HSBC is a financial firm, any risk associated with the new business line would be considered financial risk.

However, this approach is clearly too narrow, in that it excludes financial risks taken by nonfinancial firms. For example, auto manufacturers that provide financing for car buyers, large restaurant chains that hedge food prices with commodity futures, and municipalities that issues bonds to finance infrastructure projects all face financial risk.

This approach may also be too broad, in that it also includes risks to financial firms that have little to do with finance. For instance, most financial firms rely on large, complex computer systems. Should a financial risk manager try to assess the probability of network crashes, or the relative risk of two database platforms? The distribution of losses due to fires at bank branches? The risk of lawsuits arising from a new retail investment product? Lawsuits due to a new human resources policy? While a degree in finance might seem unlikely to prepare one to deal with these types of risk, in practice, the chief risk officer at a large financial firm often has a mandate which encompasses *all* types of risk. Similarly, regulators are concerned with risk to the financial system caused by financial firms, no matter where that risk comes from. Because of this, many would define financial risk management to include all aspects of financial firms, and the financial activities of nonfinancial firms. In recent years, the role of many financial risk professionals has expanded. Many welcome this increased

responsibility, while others see it as potentially dangerous mission creep. If financial risk is defined too broadly, risk managers may take responsibility for risks for which they have little or no expertise.

Another simple way to define financial risk management would be in terms of financial instruments. Defined this way, any risk arising from the use of financial instruments is within the scope of financial risk management. By this definition, the financial risk arising from the use of an interest rate swap is within the scope of financial risk management, whether the two parties involved are financial institutions or not. This is the definition preferred by many practitioners. Readers should be aware of both possibilities: that financial risk management can be defined in terms of financial firms or financial instruments.

TYPES OF FINANCIAL RISK

Financial risk is often divided into four principal types of risk: market risk, credit risk, liquidity risk, and operational risk. To varying degrees, most financial transactions involve aspects of all four types of risk. Within financial institutions, risk management groups are often organized along these lines. Because instruments with the greatest market risk tend to have the most variable liquidity risk, market risk and liquidity risk are often managed by a single group within financial firms. In addition to market risk, credit risk, liquidity risk, and operational risk, many firms will also have an enterprise risk management group, giving us a total of five principal areas of risk management. We consider each in turn.

Market Risk

Market risk is risk associated with changing asset values. Market risk is most often associated with assets that trade in liquid financial markets, such as stocks and bonds. During trading hours, the prices of stocks and bonds constantly fluctuate. An asset's price will change as new information becomes available and investors reassess the value of that asset. An asset's value can also change due to changes in supply and demand.

All financial assets have market risk. Even if an asset is not traded on an exchange, its value can change over time. Firms that use mark-to-market accounting recognize this change explicitly. For these firms, the change in value of exchange-traded assets will be based on market prices. Other assets will either be *marked to model*—that is, their prices will be determined based on financial models with inputs that may include market prices—or their prices will be based on broker quotes—that is, their prices will be based on the price at which another party expresses their willingness to buy or sell the assets. Firms that use historical cost accounting, or book value accounting, will normally only realize a profit or a loss when an asset is sold. Even if the value of the asset is not being updated on a regular basis, the asset

still has market risk. For this reason, most firms that employ historical cost accounting will reassess the value of their portfolios when they have reason to believe that there has been a significant change in the value of their assets.

For most financial instruments, we expect price changes to be relatively smooth and continuous most of the time, and large and discontinuous rarely. Because of this, market risk models often involve continuous distribution. Market risk models can also have a relatively high frequency (i.e., daily or even intraday). For many financial instruments, we will have a large amount of historical market data that we can use to evaluate market risk.

Credit Risk

Credit risk is the risk that one party in a financial transaction will fail to pay the other party. Credit risk can arise in a number of different settings. Firms may extend credit to suppliers and customers. Credit card debt and home mortgages create credit risk. One of the most common forms of credit risk is the risk that a corporation or government will fail to make interest payments or to fully repay the principal on bonds they have issued. This type of risk is known as *default risk*, and in the case of national governments it is also referred to as *sovereign risk*. Defaults occur infrequently, and the simplest models of default risk are based on discrete distributions. Although bond markets are large and credit rating agencies have been in existence for a long time, default events are rare. Because of this, we have much less historical data to work with when developing credit models, compared to market risk models.

For financial firms, counterparty credit risk is another important source of credit risk. While credit risk always involves two counterparties, when risk managers talk about counterparty credit risk they are usually talking about the risk arising from a significant long-term relationship between two counterparties. Prime brokers will often provide loans to investment firms, provide them with access to emergency credit lines, and allow them to purchase securities on margin. Assessing the credit risk of a financial firm can be difficult, time consuming, and costly. Because of this, when credit risk is involved, financial firms often enter into long-term relationships based on complex legal contracts. Counterparty risk specialists help design these contracts and play a lead role in assessing and monitoring the risk of counterparties.

Derivatives contracts can also lead to credit risk. A derivative is essentially a contract between two parties, that specifies that certain payments be made based on the value of an underlying security or securities. Derivatives include futures, forwards, swaps, and options. As the value of the underlying asset changes, so too will the value of the derivative. As the value of the derivative changes, so too will the amount of money that the counterparties owe each other. This leads to credit risk.

Another very common form of credit risk in financial markets is settlement risk. Typically, when you buy a financial asset you do not have to pay for the asset immediately. Settlement terms vary by market, but typical settlement periods are one to three days. Practitioners would describe settlement as being T+2, when payment is due two days after a trade has happened.

Liquidity Risk

Liquidity risk is the risk that you will either not be able to buy or sell an asset, or that you will not be able to buy or sell an asset in the desired quantity at the current market price. We often talk about certain markets being more or less liquid. Even in relatively liquid markets, liquidity risk can be a problem for large financial firms.

Liquidity risk can be difficult to describe mathematically, and the data needed to model liquidity risk can be difficult to obtain even under the best circumstances. Though its importance is widely recognized, liquidity risk modeling has traditionally received much less attention than market or credit risk modeling. Current approaches to liquidity risk management are often primitive. The more complex approaches that do exist are far from standard.

Operational Risk

Operational risk is risk arising from all aspects of a firm's business activities. Put simply, it is the risk that people will make mistakes and that systems will fail. Operational risk is a risk that all financial firms must deal with.

Just as the number of activities that businesses carry out is extremely large, so too are the potential sources of operational risk. That said, there are broad categories on which risk managers tend to focus. These include legal risk (most often risk arising from contracts, which may be poorly specified or misinterpreted), systems risk (risk arising from computer systems) and model risk (risk arising from pricing and risk models, which may contain errors, or may be used inappropriately).

As with credit risk, operational risk tends to be concerned with rare but significant events. Operational risk presents additional challenges in that the sources of operational risk are often difficult to identify, define, and quantify.

Enterprise Risk

The enterprise risk management group of a firm, as the name suggests, is responsible for the risk of the entire firm. At large financial firms, this often means overseeing market, credit, liquidity, and operations risk groups, and combining information from those groups into

summary reports. In addition to this aggregation role, enterprise risk management tends to look at overall business risk. Large financial companies will often have a number of business units (e.g., capital markets, corporate finance, commercial banking, retail banking, asset management, etc.). Some of these business units will work very closely with risk management (e.g. capital markets, asset management), while others may have very little day-to-day interaction with risk (e.g. corporate finance). Regardless, enterprise risk management would assess how each business unit contributes to the overall profitability of the firm in order to assess the overall risk to the firm's revenue, income, and capital.

WHAT DOES A RISK MANAGER DO?

The responsibilities of a chief risk officer (CRO) can be divided into four main tasks: defining risk, monitoring risk, controlling risk, and explaining or communicating risk. Other risk professionals will be involved in some or all of these tasks.

Defining risk is the starting point of the risk management process, and possibly the most important task. Defining risk involves clearly identifying what financial variables are to be monitored and then defining acceptable behavior for those variables. Acceptable behavior is often defined in terms of averages, minimums, and maximums. For example, we might state that net equity exposure is expected to average 10% of assets under management and will not exceed 20%, or that forecasted standard deviation of daily profits will not exceed 10% for more than one day each month and will never exceed 15%. These portfolio specifications and limits are often collected in a document detailing risk management policies and procedures. This document likely outlines who is responsible for risk management, and what action will be taken in the event that a policy is breached.

Defining risk parameters in advance helps a firm manage its investments in a consistent and transparent manner. If done correctly a well-defined risk framework will make the investment process more predictable and help reduce extrinsic risk. For example, most hedge funds are allowed to invest in a wide range of financial products and to use considerable leverage. If there were no risk limits, risk levels could vary widely. By carefully defining how risk is going to be managed and communicating this to investors, we can significantly reduce extrinsic risk.

It is worth pointing out that the job of a risk manager is not necessarily to reduce risk. For an investment firm, more risk is often associated with higher potential profits. An investor might be just as worried about risk being too low as too high.

Sophisticated investors can adjust their level of risk by increasing or decreasing their exposure to a fund or by hedging. In order to do this, they need as much information as possible about the risks that the fund is taking. The risk manager can reduce extrinsic risk for these

investors and help them achieve a more optimal allocation, by accurately communicating the risks that their fund is taking.

After we have defined the risk parameters of a portfolio, we need to monitor these parameters. This is the task that is most frequently associated with the role of risk management. You can imagine a CRO striding into the chief investment officer's office, to report that the firm's expected standard deviation has increased recently and is getting very close to its limit. Monitoring risk in a timely manner can often be technologically challenging.

The third, and possibly most important, task for a risk manager is to control or manage risk. Risk can be managed in a number of ways. As well as helping to enforce limits, at some investment firms the CRO will actually manage, or help manage, a hedge portfolio, which is used to control risk. At other firms, risk managers will work more closely with portfolio managers, adjusting the portfolio as necessary to increase or decrease risk.

In addition to communicating with their colleagues (e.g., back office personnel, traders, portfolio managers), an increasingly important job for risk managers is communicating with regulators and investors. In all cases, the risk manager is engaged in what we might call dimensionality reduction, taking a large set of financial instruments and market data and reducing them to small number of key statistics and insights.

A VERY BRIEF HISTORY OF RISK MANAGEMENT

Christiaan Huygens was a Dutch polymath whose interest ranged from astronomy to mathematics to engineering. Among other accomplishments Huygens discovered Titan, the largest moon of Saturn, and helped design the water fountains at the Palace of Versailles outside of Paris. But it was Huygens's publication of *De Ratiociniis in Ludo Aleae*, or *On Reason in Games of Chance*, in 1657 that is of importance to the study of risk. It was in this book that Huygens first developed the formal concept of expectations.

Like many of his contemporaries, Huygens was interested in games of chance. As he described it, if a game has a 50% probability of paying \$3 and a 50% probability of paying \$7, then this is, in a way, equivalent to having \$5 with certainty. This is because we expect, on average, to win \$5 in this game:

$$50\% \times \$3 + 50\% \times \$7 = \$5 \quad (1.1)$$

We'll have a lot more to say about expectations in Chapter 2.

As early as 1713, Daniel and Nicolas Bernoulli were beginning to doubt that human beings were quite so logical when it came to evaluating risks, and, as we will see later when we explore behavioral finance, economist still struggle with this topic. Beyond the evaluation of games of chance, the more general concept of expectations is the basis for our modern

definitions for mean, variance, and many other statistical concepts. It is arguably the most important concept in modern statistics.

From the 18th century, we jump to the Crash of 1929. Even after the financial crisis of 2008 and the ensuing Great Recession, the Crash of 1929 is still considered by most experts to have been the worst financial crash in history. On October 28, 1929, Black Monday, the Dow Jones Industrial Average lost 13%. For the entire month of October the index was down 20%. The crash was likely a leading cause of the ensuing Great Depression. That the crash of the financial markets could have such a profound impact on the rest of the economy was a clear indication of the central role that financial markets play in modern economies. This potential for widespread harm is a major justification for financial regulation. In the wake of the crash, the United States government passed the Securities Act of 1933 and the Securities Exchange Act of 1934. The former would, among other things, go on to become the defining regulation for hedge funds in the United States. The latter established the Securities and Exchange Commission (SEC). More recent regulatory efforts, including the Basel Accords, are direct descendants of these efforts. Today, for better or worse, regulatory compliance is a full-time job for many financial risk managers.

In 1952 *The Journal of Finance* published “Portfolio Selection” by Harry Markowitz. The article introduced the world to Modern Portfolio Theory (MPT). For this and related work, Markowitz would go on to win the Nobel Prize in Economics. The key insight of MPT is that investors are trying to get the highest returns with the least amount of risk. Given two portfolios with the same level of risk but different expected returns, a rational investor will prefer the portfolio with the higher expected return. Similarly, given two portfolios with the same expected return, but different risk levels, a rational investor will prefer the less risky portfolio. That this seems obvious—that it seems natural to frame investing in terms of risk and return—is a testament to the profound impact of MPT on finance and risk management. As mentioned previously, a risk manager’s job is not necessarily to reduce risk. If we reduce risk but also reduce returns, investors may not be better off.

In his initial paper, Markowitz modeled risk in terms of variance or standard deviation. Standard deviation is still one of the most widely used measures for characterizing risk. As we will see in the next chapter, though, risk management has moved far beyond this narrow definition of risk.


On Monday October 19, 1987, stock markets around the world crashed. The Dow Jones Industrial Average lost 22%, and the S&P 500 lost 20%. This was the worst recorded one-day return in the history of both indexes. Today, when people talk about Black Monday, more often than not they are referring to this event and not the previous Black Monday from 1929. Oddly, this more recent Black Monday was a relatively isolated incident. The S&P 500 was actually up for 1987, and the economy grew both in 1987 and 1988. Contrast

this to 1929, where the economy shrank by 26% over the three succeeding years. To this day, the causes of Black Monday '87 are still debated. The growing use of program trading and portfolio insurance were both possible causes. Whether or not they caused the crash, they did cause an increase in trading volume. This already high trading volume spiked on Black Monday, and many markets were unable to cope. As a result of Black Monday many exchanges upgraded their trading systems. Most financial firms that use stress testing have at least one stress test based on Black Monday '87. We will have more to say about stress testing in Chapter 4.

Long-Term Capital Management (LTCM) belongs in the history of risk management for many reasons. The failure of LTCM in 1998 was a shock to the financial community. At the time it was one of the largest, most highly regarded hedge funds in the world, with an enviable track record. It could count among its founders John Meriwether, Myron Scholes, and Robert Merton. LTCM's portfolio was so large and its network of counterparties so extensive that regulators feared its failure could devastate the larger financial community. The New York Fed called an emergency meeting of the heads of the largest investment banks to intercede and liquidate LTCM's portfolio in an orderly fashion. That a firm virtually unknown outside of the financial community, employing fewer than 200 people, could pose a risk to the entire financial system highlighted the incredible growth that had taken place throughout the 1980s and 1990s outside of traditional financial markets, most notably in derivatives and hedge funds. This is the first reason that justifies LTCM's place in the history of risk management. The second reason, which is often overlooked in the story of LTCM, is that its founders viewed risk management as a means to reduce capital requirements for both financial and nonfinancial firms. In theory, by freeing up capital for use in other endeavors, firms and the economy as a whole could become more productive. In retrospect, LTCM was clearly overleveraged and undercapitalized, but the idea that risk management can create real efficiencies for firms is still sound. The final—and perhaps most obvious—reason that LTCM belongs in the history of risk management is that its failure was in part due to its risk models and how those models were used. The risk management community took notice. The external event that precipitated LTCM's failure, Russia's default on its domestic debt, is still the basis of stress tests at many financial firms. More importantly it highlighted the limits of historically based quantitative risk models and spurred the risk management community to look for more robust solutions.

RiskMetrics is another important organization in risk management history. From its spin-off from JP Morgan in 1998 to its IPO in 2008 to its eventual acquisition by MSCI Inc. in 2010, RiskMetrics was one of the largest and most successful companies devoted entirely to risk management. Figure 1.1, which shows RiskMetrics' annual revenue, gives

FIGURE 1.1 RiskMetrics Annual Revenue




Source: Securities and Exchange Commission

an indication of just how quickly the firm grew. RiskMetrics software is still used to help manage the risk of some of the largest financial firms in the world. Despite these impressive facts, more than anything else, RiskMetrics earned its place in the history of risk management for what it gave away for free. In 1992 RiskMetrics published the *RiskMetrics Technical Document*. The document outlined RiskMetrics’s approach to evaluating risk, and popularized the concept of value at risk, or VaR. Along with standard deviation, VaR is one of the most popular statistics for summarizing financial risk.

The Global Association of Risk Professionals (GARP), the world’s largest nonprofit financial risk management association, was founded in 1996. As of 2018 GARP has over 150,000 members in 195 countries. In addition to hosting conferences, producing publications, and providing continuing education, GARP is the sponsor of the Financial Risk Manager (FRM) Exam.¹ The extraordinary growth of GARP’s membership and of the number of people taking the FRM exam annually is an indication not only of the growth of the risk management industry, but also of the increasing importance of standards within the industry. Figure 1.2 shows the increase in FRM Exam enrollment over time. The exam was changed from a one-year to a two-year format starting in 2009, leading to a temporary spike in enrollment that year.

¹Full disclosure: The author is a longstanding member of GARP and an FRM holder.

FIGURE 1.2 FRM Exam Enrollment

Source: Global Association of Risk Professionals

THE FUTURE OF RISK MANAGEMENT

Despite tremendous growth in recent years, financial risk management is still a young discipline. We can expect to see many changes in the roles of financial risk managers in the coming years.

The financial crisis of 2008 called into question some quantitative risk models, but it also caused many to argue for a greater role for risk managers within financial firms. While there were certainly instances when models were used incorrectly, the far greater problem was that the decision makers at large financial firms either never received the data they needed, didn't understand it, or chose to ignore it. It was not so much that we lacked the tools to properly assess risk, as it was that the tools were not being used or that the people using the tools were not being listened to.

As risk management continues to gain wider acceptance, the role of risk managers in communicating with investors and regulators will continue to grow. We are also likely to see an increasingly integrated approach to risk management and performance analysis, which are now treated as separate activities by most financial firms.

There are still important areas of risk management, such as liquidity risk, where widely accepted models and standards have yet to be developed. If history is any guide, financial markets will continue to grow in breadth, speed, and complexity. Along with this growth will come new challenges for risk managers.

The future of risk management is very bright.

2

MARKET RISK: STANDARD DEVIATION

In this chapter we will start our exploration of market risk. For the time being, we will ignore all other sources of risk, including liquidity and credit risk. We begin by examining the market risk of one security.

This is the first of six chapters on market risk. As we will see, risk managers have many statistics to choose from when trying to measure market risk. No single statistic is perfect. Each has its strengths and weaknesses. Portfolio managers, investors, and regulators depend on risk managers to choose the right statistics to report. Over the course of the next six chapters, pay special attention to: what assumptions are being made, how the various statistics compare to each other, which aspects of market risk they describe well, and which aspects of market risk they omit.

RISK AND STANDARD DEVIATION

The classical risk model equates risk solely with the standard deviation of returns. As mentioned in the previous chapter, Markowitz's market portfolio theory (MPT) assumes that investors equate standard deviation with risk. Today's risk managers have not abandoned standard deviation; rather most report standard deviation along with additional statistics, such as skewness, kurtosis, and value at risk. We will explore these other statistics eventually, but we begin with a review of standard deviation.

Standard deviation is so widely used throughout finance that we often refer to it as *volatility*, or simply *vol*. While it is important to be aware of this practice, standard deviation

is a precise mathematical term, whereas volatility is a more general concept. Except for in rare instances, when we are actually referring to this more general concept of volatility, we will use *standard deviation* throughout the rest of the text.

Derivatives present a unique problem for risk management. The payout profiles of options are typically discontinuous. Small changes in the return of the underlying security have the potential to create large changes in the value of the derivative. Because of this, derivatives are often described as being nonlinear. We will explore methods for dealing with options in subsequent chapters. For the remainder of this chapter, when we refer to *security returns*, assume that we are referring to instruments such as equities, floating exchange rates, bonds, or commodity futures. Unless noted otherwise, assume that the securities are liquid, and that there is no credit risk.

Before we can formally define standard deviation, we will need to define the mean of a random variable. To this end, we begin by discussing averages, random variables, and expectations. It is assumed that the reader is familiar with probabilities, random variables, and probability density functions.¹

AVERAGES

Everybody knows what an average is. We come across averages every day, whether they are earned run averages in baseball or grade point averages in school. In statistics there are actually three different types of averages: means, modes, and medians. By far the most commonly used in risk management is the mean.

Population and Sample Data

If you wanted to know the mean age of people working in your firm, you would simply ask every person in the firm his or her age, add the ages together, and divide by the number of people in the firm. Assuming there are n employees and a_i is the age of the i th employee, then the mean, μ , is simply

$$\mu = \frac{1}{n} \sum_{i=1}^n a_i = \frac{1}{n} (a_1 + a_2 + \dots + a_{n-1} + a_n) \quad (2.1)$$

It is important for us to differentiate between population statistics and sample statistics. In this example, μ is the population mean. Assuming nobody lied about his or her age, and forgetting about rounding errors and other trivial details, we know the mean age of the

¹For readers not familiar with these topics, an overview can be found in Chapter 2 of *Mathematics and Statistics for Financial Risk Management*, Miller (2014).

people in your firm *exactly*. We have a complete data set of everybody in your firm; we've surveyed the entire population.

This state of absolute certainty is, unfortunately, quite rare in finance. More often, we are faced with a situation such as this: Estimate the mean return of stock ABC, given the most recent year of daily returns. In a situation like this, we assume there is some underlying data-generating process with statistical properties that are constant over time. The underlying process has a true mean, but we cannot observe it directly. We can only estimate the true mean based on our limited data sample. In our example, assuming n returns, we estimate the mean using the same formula as before

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n r_i = \frac{1}{n}(r_1 + r_2 + \dots + r_{n-1} + r_n) \quad (2.2)$$

where $\hat{\mu}$ (pronounced “mu hat”) is our *estimate* of the true mean, μ , based on our sample of n returns. We call this the sample mean.

The median and mode are also types of averages. They are used less frequently in finance, but both can be useful. The median represents the center of a group of data; within the group, half the data points will be less than the median and half will be greater. The mode is the value that occurs most frequently.

SAMPLE PROBLEM

Question:

Calculate the mean, median, and mode of the following data set:

−20%, −10%, −5%, −5%, 0%, 10%, 10%, 10%, 19%

Answer:

$$\text{Mean} = \frac{1}{9}(-20\% - 10\% - 5\% - 5\% + 0\% + 10\% + 10\% + 10\% + 19\%) = 1\%$$

$$\text{Mode} = 10\%$$

$$\text{Median} = 0\%$$

If there is an even number of data points, the median is found by averaging the two centermost points. In the following series:

5%, 10%, 20%, 25%

the median is 15%. The median can be useful for summarizing data that is asymmetrical or contains significant outliers. Financial time series are often asymmetrical and often contain

outliers. Both Black Monday events, which we mentioned in the last chapter, are examples of significant outliers.

A data set can also have more than one mode. If the maximum frequency is shared by two or more values, all of those values are considered modes. In the following series, the modes are 10% and 20%:

5%, 10%, 10%, 10%, 14%, 16%, 20%, 20%, 20%, 24%

In calculating the mean in Equations 2.1 and 2.2, each data point was counted exactly once. In certain situations, we might want to give more or less weight to certain data points. In calculating the average return of stocks in an equity index, we might want to give more weight to larger firms, perhaps weighting their returns in proportion to their market capitalizations. Given n data points, $x_i = x_1, x_2, \dots, x_n$, with corresponding weights, w_i , we can define the weighted mean, μ_w , as

$$\mu_w = \frac{\sum_{i=1}^n w_i x_i}{\sum_{i=1}^n w_i} \quad (2.3)$$

The standard mean from Equation 2.1 can be viewed as a special case of the weighted mean, where all the values have equal weight.

Discrete Random Variables

Financial markets are highly uncertain. Day-to-day changes to the price of a stock, a foreign exchange rate, or interest rates, are, in most cases, essentially random. In statistics we can model random, or *stochastic*, phenomena using random variables. There are two basic types of random variables: discrete and continuous. A discrete random variable can take on only a finite number of possible values. For example, a bond that is worth either \$60 if it does default, or \$100 if it does not default. A continuous random variable can take on any value within a given range. That range can be finite or infinite. For example, for a stock index, we might assume that the returns can be any value from -100% to infinity.

For a discrete random variable, we can also calculate the mean, median, and mode. For a random variable, X , with possible values, x_i , and corresponding probabilities, p_i , we define the mean, μ , as

$$\mu = \sum_{i=1}^n p_i x_i \quad (2.4)$$

The equation for the mean of a discrete random variable is a special case of the weighted mean, where the outcomes are weighted by their probabilities, and the sum of the weights is equal to one.