

WILEY FINANCE

Mathematics and Statistics for Financial Risk Management

Michael B. Miller

Mathematics and Statistics for 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.

Mathematics and Statistics for Financial Risk Management

MICHAEL B. MILLER



WILEY

John Wiley & Sons, Inc.

Copyright © 2012 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 <http://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 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.

Library of Congress Cataloging-in-Publication Data:

Miller, Michael B.

Mathematics and statistics for financial risk management / Michael B. Miller.

p. cm. – (Wiley finance)

Includes bibliographical references and index.

ISBN 978-1-118-17062-5 (hardback); ISBN 978-1-118-22777-0 (ebk);

ISBN 978-1-118-24419-7 (ebk); ISBN 978-1-118-23976-6 (ebk)

1. Risk management—Mathematical models. 2. Risk management—Statistical methods.

I. Title.

HD61.M537 2012

332.01'5195—dc23

2011039256

ISBN 978-1-118-17062-5

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Contents

Preface	ix
Acknowledgments	xi
CHAPTER 1	
Some Basic Math	1
Logarithms	1
Log Returns	3
Compounding	4
Limited Liability	5
Graphing Log Returns	5
Continuously Compounded Returns	7
Combinatorics	9
Discount Factors	10
Geometric Series	11
Problems	16
CHAPTER 2	
Probabilities	19
Discrete Random Variables	19
Continuous Random Variables	20
Mutually Exclusive Events	26
Independent Events	27
Probability Matrices	28
Conditional Probability	30
Bayes' Theorem	31
Problems	36
CHAPTER 3	
Basic Statistics	39
Averages	39
Expectations	46
Variance and Standard Deviation	51
Standardized Variables	54

Covariance	54
Correlation	56
Application: Portfolio Variance and Hedging	57
Moments	60
Skewness	60
Kurtosis	64
Coskewness and Cokurtosis	67
Best Linear Unbiased Estimator (BLUE)	71
Problems	72
CHAPTER 4	
Distributions	75
Parametric Distributions	75
Uniform Distribution	75
Bernoulli Distribution	78
Binomial Distribution	79
Poisson Distribution	83
Normal Distribution	84
Lognormal Distribution	88
Central Limit Theorem	90
Application: Monte Carlo Simulations Part I: Creating Normal Random Variables	92
Chi-Squared Distribution	94
Student's t Distribution	95
F-Distribution	97
Mixture Distributions	99
Problems	102
CHAPTER 5	
Hypothesis Testing & Confidence Intervals	105
The Sample Mean Revisited	105
Sample Variance Revisited	107
Confidence Intervals	108
Hypothesis Testing	109
Chebyshev's Inequality	113
Application: VaR	114
Problems	124
CHAPTER 6	
Matrix Algebra	127
Matrix Notation	127
Matrix Operations	129

Application: Transition Matrices	136
Application: Monte Carlo Simulations Part II: Cholesky Decomposition	138
Problems	141
CHAPTER 7	
Vector Spaces	143
Vectors Revisited	143
Orthogonality	146
Rotation	152
Principal Component Analysis	157
Application: The Dynamic Term Structure of Interest Rates	162
Application: The Structure of Global Equity Markets	167
Problems	171
CHAPTER 8	
Linear Regression Analysis	173
Linear Regression (One Regressor)	173
Linear Regression (Multivariate)	183
Application: Factor Analysis	188
Application: Stress Testing	192
Problems	194
CHAPTER 9	
Time Series Models	197
Random Walks	197
Drift-Diffusion	199
Autoregression	200
Variance and Autocorrelation	205
Stationarity	206
Moving Average	212
Continuous Models	212
Application: GARCH	215
Application: Jump-Diffusion	217
Application: Interest Rate Models	218
Problems	220
CHAPTER 10	
Decay Factors	223
Mean	223
Variance	230
Weighted Least Squares	231

Other Possibilities	232
Application: Hybrid VaR	233
Problems	234
APPENDIX A	
Binary Numbers	237
APPENDIX B	
Taylor Expansions	239
APPENDIX C	
Vector Spaces	241
APPENDIX D	
Greek Alphabet	242
APPENDIX E	
Common Abbreviations	243
Answers	245
References	283
About the Author	285
Index	287


Preface

The recent financial crisis and its impact on the broader economy underscore the importance of financial risk management in today's world. At the same time, financial products and investment strategies are becoming increasingly complex. It is more important than ever that risk managers possess a sound understanding of mathematics and statistics.

Mathematics and Statistics for Financial Risk Management is a guide to modern financial risk management for both practitioners and academics. Risk management has made great strides in recent years. Many of the mathematical and statistical tools used in risk management today were originally adapted from other fields. As the field has matured, risk managers have refined these tools and developed their own vocabulary for characterizing risk. As the field continues to mature, these tools and vocabulary are becoming increasingly standardized. By focusing on the application of mathematics and statistics to actual risk management problems, this book helps bridge the gap between mathematics and statistics in theory and risk management in practice.

Each chapter in this book introduces a different topic in mathematics or statistics. As different techniques are introduced, sample problems and application sections demonstrate how these techniques can be applied to actual risk management problems. Exercises at the end of each chapter, and the accompanying solutions at the end of the book, allow readers to practice the techniques they are learning and to monitor their progress.

This book assumes that readers have a solid grasp of algebra and at least a basic understanding of calculus. Even though most chapters start out at a very basic level, the pace is necessarily fast. For those who are already familiar with the topic, the beginning of each chapter will serve as a quick review and as an introduction to certain vocabulary terms and conventions. Readers who are new to these topics may find they need to spend more time in the initial sections.

Risk management in practice often requires building models using spreadsheets or other financial software. Many of the topics in this book are accompanied by a  icon, as shown here. These icons indicate that Excel examples can be found at John Wiley & Sons' companion website for

Mathematics and Statistics for Financial Risk Management, at www.wiley.com/go/millerfinance.

You can also visit the author's web site, www.risk256.com, for the latest financial risk management articles, code samples, and more. To provide feedback, you can contact the author at mike@risk256.com.

Acknowledgments

Like most of today's risk managers, I learned much of what I know about risk management on the job. I was fortunate to work with some very knowledgeable individuals early in my career. In particular, I would like to thank Gideon Pell and Kent Osband.

This book would not have been possible without the help of many individuals. I would like to thank Jeffrey Garnett, Steve Lerit, Hyunsuk Moon, Elliot Noma, and Eldar Radovici for taking time to read early drafts. The book is certainly better for their comments and feedback. I would also like to thank everybody at John Wiley & Sons for their help in bringing this book together.

Finally, and most importantly, I would like to thank my wife, Amy, who not only read over early drafts and talked me through a number of decisions, but also put up with countless nights and weekends of typing and editing. For this and much, much more, thank you.

Some Basic Math

In this chapter we will review three math topics—logarithms, combinatorics, and geometric series—and one financial topic, discount factors. Emphasis will be given to the specific aspects of these topics that are most relevant to risk management.

LOGARITHMS

In mathematics, logarithms, or logs, are related to exponents, as follows:

$$\log_b a = x \Leftrightarrow a = b^x \quad (1.1)$$

We say, “The log of a , base b , equals x , which implies that a equals b to the x and vice versa.” If we take the log of the right-hand side of Equation 1.1 and use the identity from the left-hand side of the equation, we can show that:

$$\log_b(b^x) = x \quad (1.2)$$

Taking the log of b^x effectively cancels out the exponentiation, leaving us with x .

An important property of logarithms is that the logarithm of the product of two variables is equal to the sum of the logarithms of those two variables. For two variables, X and Y :

$$\log_b(XY) = \log_b X + \log_b Y \quad (1.3)$$

Similarly, the logarithm of the ratio of two variables is equal to the difference of their logarithms:

$$\log_b\left(\frac{X}{Y}\right) = \log_b X - \log_b Y \quad (1.4)$$

If we replace Y with X in Equation 1.3, we get:

$$\log_b(X^2) = 2 \log_b X \quad (1.5)$$

We can generalize this result to get the following power rule:

$$\log_b(X^n) = n \log_b X \quad (1.6)$$

In general, the base of the logarithm, b , can have any value. Base 10 and base 2 are popular bases in certain fields, but in many fields, and especially in finance, e , Euler's number, is by far the most popular. Base e is so popular that mathematicians have given it its own name and notation. When the base of a logarithm is e , we refer to it as a *natural logarithm*. In formulas, we write:

$$\ln(a) = x \Leftrightarrow a = e^x \quad (1.7)$$

From this point on, unless noted otherwise, assume that any mention of logarithms refers to natural logarithms.

Logarithms are defined for all real numbers greater than or equal to zero. Figure 1.1 shows a plot of the logarithm function. The logarithm of zero is negative infinity, and the logarithm of one is zero. The function




FIGURE 1.1 Natural Logarithm

grows without bound; that is, as X approaches infinity, the $\ln(X)$ approaches infinity as well.

LOG RETURNS

One of the most common applications of logarithms in finance is computing log returns. Log returns are defined as follows:

$$r_t \equiv \ln(1 + R_t) \quad \text{where} \quad R_t = \frac{P_t - P_{t-1}}{P_{t-1}} \quad (1.8)$$

Here r_t is the log return at time t , R_t is the standard or simple return, and P_t is the price of the security at time t . We use this convention of capital R for simple returns and lowercase r for log returns throughout the rest of the book. This convention is popular, but by no means universal. Also, be careful: Despite the name, the log return is not the log of R_t , but the log of $(1 + R_t)$.

For small values, log returns and simple returns will be very close in size. A simple return of 0% translates exactly to a log return of 0%. A simple return of 10% translates to a log return of 9.53%. That the values are so close is convenient for checking data and preventing operational errors. Table 1.1 shows some additional simple returns along with their corresponding log returns.

TABLE 1.1 Log Returns and Simple Returns

R	$\ln(1 + R)$
1.00%	1.00%
5.00%	4.88%
10.00%	9.53%
20.00%	18.23%

To get a more precise estimate of the relationship between standard returns and log returns, we can use the following approximation:*

$$r \approx R - \frac{1}{2}R^2 \quad (1.9)$$

*This approximation can be derived by taking the Taylor expansion of Equation 1.8 around zero. Though we have not yet covered the topic, for the interested reader a brief review of Taylor expansions can be found in Appendix B.

As long as R is small, the second term on the right-hand side of Equation 1.9 will be negligible, and the log return and the simple return will have very similar values.

COMPOUNDING

Log returns might seem more complex than simple returns, but they have a number of advantages over simple returns in financial applications. One of the most useful features of log returns has to do with compounding returns. To get the return of a security for two periods using simple returns, we have to do something that is not very intuitive, namely adding one to each of the returns, multiplying, and then subtracting one:

$$R_{2,t} = \frac{P_t - P_{t-2}}{P_{t-2}} = (1 + R_{1,t})(1 + R_{1,t-1}) - 1 \quad (1.10)$$

Here the first subscript on R denotes the length of the return, and the second subscript is the traditional time subscript. With log returns, calculating multiperiod returns is much simpler; we simply add:

$$r_{2,t} = r_{1,t} + r_{1,t-1} \quad (1.11)$$

By substituting Equation 1.8 into Equation 1.10 and Equation 1.11, you can see that these are equivalent. It is also fairly straightforward to generalize this notation to any return length.

SAMPLE PROBLEM

Question:

Using Equation 1.8 and Equation 1.10, generalize Equation 1.11 to returns of any length.

Answer:

$$R_{n,t} = \frac{P_t - P_{t-n}}{P_{t-n}} = \frac{P_t}{P_{t-n}} - 1 = \frac{P_t}{P_{t-1}} \frac{P_{t-1}}{P_{t-2}} \cdots \frac{P_{t-n+1}}{P_{t-n}} - 1$$

$$R_{n,t} = (1 + R_{1,t})(1 + R_{1,t-1}) \cdots (1 + R_{1,t-n+1}) - 1$$

$$(1 + R_{n,t}) = (1 + R_{1,t})(1 + R_{1,t-1}) \cdots (1 + R_{1,t-n+1})$$

$$r_{n,t} = r_{1,t} + r_{1,t-1} + \cdots + r_{1,t-n+1}$$

Note that to get to the last line, we took the logs of both sides of the previous equation, using the fact that the log of the product of any two variables is equal to the sum of their logs, as shown in Equation 1.3.

LIMITED LIABILITY

Another useful feature of log returns relates to limited liability. For many financial assets, including equities and bonds, the most that you can lose is the amount that you've put into them. For example, if you purchase a share of XYZ Corporation for \$100, the most you can lose is that \$100. This is known as limited liability. Today, limited liability is such a common feature of financial instruments that it is easy to take it for granted, but this was not always the case. Indeed, the widespread adoption of limited liability in the nineteenth century made possible the large publicly traded companies that are so important to our modern economy, and the vast financial markets that accompany them.

That you can lose only your initial investment is equivalent to saying that the minimum possible return on your investment is -100% . At the other end of the spectrum, there is no upper limit to the amount you can make in an investment. The maximum possible return is, in theory, infinite. This range for simple returns, -100% to infinity, translates to a range of negative infinity to positive infinity for log returns.

$$\begin{aligned}R_{\min} = -100\% &\Rightarrow r_{\min} = -\infty \\R_{\max} = +\infty &\Rightarrow r_{\max} = +\infty\end{aligned}\tag{1.12}$$

As we will see in the following chapters, when it comes to mathematical and computer models in finance, it is often much easier to work with variables that are unbounded, that is variables that can range from negative infinity to positive infinity.

GRAPHING LOG RETURNS

Another useful feature of log returns is how they relate to log prices. By rearranging Equation 1.10 and taking logs, it is easy to see that:

$$r_t = p_t - p_{t-1}\tag{1.13}$$

where p_t is the log of P_t , the price at time t . To calculate log returns, rather than taking the log of one plus the simple return, we can simply calculate the logs of the prices and subtract.

Logarithms are also useful for charting time series that grow exponentially. Many computer applications allow you to chart data on a logarithmic scale. For an asset whose price grows exponentially, a logarithmic scale prevents the compression of data at low levels. Also, by rearranging Equation 1.13, we can easily see that the change in the log price over time is equal to the log return:

$$\Delta p_t = p_t - p_{t-1} = r_t \quad (1.14)$$

It follows that, for an asset whose return is constant, the change in the log price will also be constant over time. On a chart, this constant rate of change over time will translate into a constant slope. Figures 1.2 and 1.3 both show an asset whose price is increasing by 20% each year. The y-axis for the first chart shows the price; the y-axis for the second chart displays the log price.

For the chart in Figure 1.2, it is hard to tell if the rate of return is increasing or decreasing over time. For the chart in Figure 1.3, the fact that




FIGURE 1.2 Normal Prices




FIGURE 1.3 Log Prices

the line is straight is equivalent to saying that the line has a constant slope. From Equation 1.14 we know that this constant slope is equivalent to a constant rate of return.

In the first chart, the y-axis could just have easily been the actual price (on a log scale), but having the log prices allows us to do something else. Using Equation 1.13, we can easily estimate the log return. Over 10 periods, the log price increases from approximately 4.6 to 6.4. Subtracting and dividing gives us $(6.4 - 4.6)/10 = 18\%$. So the log return is 18% per period, which—because log returns and simple returns are very close for small values—is very close to the actual simple return of 20%.

CONTINUOUSLY COMPOUNDED RETURNS

Another topic related to the idea of log returns is continuously compounded returns. For many financial products, including bonds, mortgages, and credit cards, interest rates are often quoted on an annualized periodic or nominal basis. At each payment date, the amount to be paid is equal to this nominal rate, divided by the number of periods, multiplied by some notional amount.

For example, a bond with monthly coupon payments, a nominal rate of 6%, and a notional value of \$1,000, would pay a coupon of \$5 each month: $(6\% \times \$1,000)/12 = \5 .

How do we compare two instruments with different payment frequencies? Are you better off paying 5% on an annual basis or 4.5% on a monthly basis? One solution is to turn the nominal rate into an annualized rate:

$$R_{\text{Annual}} = \left(1 + \frac{R_{\text{Nominal}}}{n}\right)^n - 1 \quad (1.15)$$

where n is the number of periods per year for the instrument.

If we hold R_{Annual} constant as n increases, R_{Nominal} gets smaller, but at a decreasing rate. Though the proof is omitted here, using L'Hôpital's rule, we can prove that, at the limit, as n approaches infinity, R_{Nominal} converges to the log rate. As n approaches infinity, it is as if the instrument is making infinitesimal payments on a continuous basis. Because of this, when used to define interest rates the log rate is often referred to as the continuously compounded rate, or simply the continuous rate. We can also compare two financial products with different payment periods by comparing their continuous rates.

SAMPLE PROBLEM

Question:

You are presented with two bonds. The first has a nominal rate of 20% paid on a semiannual basis. The second has a nominal rate of 19% paid on a monthly basis. Calculate the equivalent continuously compounded rate for each bond. Assuming both bonds have the same credit quality and are the same in all other respects, which is the better investment?

Answer:

First we compute the annual yield for both bonds:

$$R_{1,\text{Annual}} = \left(1 + \frac{20\%}{2}\right)^2 - 1 = 21.00\%$$

$$R_{2,\text{Annual}} = \left(1 + \frac{19\%}{12}\right)^{12} - 1 = 20.75\%$$

Next we convert these annualized returns into continuously compounded returns:

$$r_1 = \ln(1 + R_{1,\text{Annual}}) = 19.06\%$$

$$r_2 = \ln(1 + R_{2,\text{Annual}}) = 18.85\%$$

All other things being equal, the first bond is a better investment. We could base this on a comparison of either the annual or the continuously compounded rates.

COMBINATORICS

In elementary combinatorics, one typically learns about combinations and permutations. Combinations tell us how many ways we can arrange a number of objects, regardless of the order, whereas permutations tell us how many ways we can arrange a number of objects, taking into account the order.

As an example, assume we have three hedge funds, denoted X, Y, and Z. We want to invest in two of the funds. How many different ways can we invest? We can invest in X and Y, X and Z, or Y and Z. That's it.

In general, if we have n objects and we want to choose k of those objects, the number of combinations, $C(n, k)$, can be expressed as:

$$C(n, k) = \binom{n}{k} = \frac{n!}{k!(n-k)!} \quad (1.16)$$

where $n!$ is n factorial, such that:

$$n! = \begin{cases} 1 & n = 0 \\ n(n-1)(n-2) \cdots 1 & n > 0 \end{cases} \quad (1.17)$$

In our example with the three hedge funds, we would substitute $n = 3$ and $k = 2$, to get three possible combinations.

What if the order mattered? What if instead of just choosing two funds, we needed to choose a first-place fund and a second-place fund? How many

ways could we do that? The answer is the number of permutations, which we express as:

$$P(n, k) = \frac{n!}{(n - k)!} \quad (1.18)$$

For each combination, there are $k!$ ways in which the elements of that combination can be arranged. In our example, each time we choose two funds, there are two ways that we can order them, so we would expect twice as many permutations. This is indeed the case. Substituting $n = 3$ and $k = 2$ into Equation 1.18, we get six permutations, which is twice the number of combinations computed previously.

Combinations arise in a number of risk management applications. The binomial distribution, which we will introduce in Chapter 4, is defined using combinations. The binomial distribution, in turn, can be used to model defaults in simple bond portfolios or to back-test Value at Risk (VaR) models, as we will see in Chapter 5.

Combinations are also central to the binomial theorem. Given two variables, x and y , and a positive integer, n , the binomial theorem states:

$$(x + y)^n = \sum_{k=0}^n \binom{n}{k} x^{n-k} y^k \quad (1.19)$$

For example:

$$(x + y)^3 = x^3 + 3x^2y + 3xy^2 + y^3 \quad (1.20)$$

The binomial theorem can be useful when computing statistics such as variance, skewness, and kurtosis, which will be discussed in Chapter 3.

DISCOUNT FACTORS

Most people have a preference for present income over future income. They would rather have a dollar today than a dollar one year from now. This is why banks charge interest on loans, and why investors expect positive returns on their investments. Even in the absence of inflation, a rational person should prefer a dollar today to a dollar tomorrow. Looked at another way, we should require more than one dollar in the future to replace one dollar today.

In finance we often talk of discounting cash flows or future values. If we are discounting at a fixed rate, R , then the present value and future value are related as follows:

$$V_t = \frac{V_{t+n}}{(1 + R)^n} \quad (1.21)$$

where V_t is the value of the asset at time t and V_{t+n} is the value of the asset at time $t + n$. Because R is positive, V_t will necessarily be less than V_{t+n} . All else being equal, a higher discount rate will lead to a lower present value. Similarly, if the cash flow is further in the future—that is, n is greater—then the present value will also be lower.

Rather than work with the discount rate, R , it is sometimes easier to work with a discount factor. In order to obtain the present value, we simply multiply the future value by the discount factor:

$$V_t = \left(\frac{1}{1 + R} \right)^n V_{t+n} = \delta^n V_{t+n} \quad (1.22)$$

Because δ is less than one, V_t will necessarily be less than V_{t+n} . Different authors refer to δ or δ^n as the discount factor. The concept is the same, and which convention to use should be clear from the context.

GEOMETRIC SERIES

In the following two subsections we introduce geometric series. We start with series of infinite length. It may seem counterintuitive, but it is often easier to work with series of infinite length. With results in hand, we then move on to series of finite length in the second subsection.

Infinite Series

The ancient Greek philosopher Zeno, in one of his famous paradoxes, tried to prove that motion was an illusion. He reasoned that, in order to get anywhere, you first had to travel half the distance to your ultimate destination. Once you made it to the halfway point, though, you would still have to travel half the remaining distance. No matter how many of these half journeys you completed, there would always be another half journey left. You could never possibly reach your destination.

While Zeno's reasoning turned out to be wrong, he was wrong in a very profound way. The infinitely decreasing distances that Zeno struggled with

foreshadowed calculus, with its concept of change on an infinitesimal scale. Also, an infinite series of a variety of types turn up in any number of fields. In finance, we are often faced with series that can be treated as infinite. Even when the series is long, but clearly finite, the same basic tools that we develop to handle infinite series can be deployed.

In the case of the original paradox, we are basically trying to calculate the following summation:

$$S = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots \quad (1.23)$$

What is S equal to? If we tried the brute force approach, adding up all the terms, we would literally be working on the problem forever. Luckily, there is an easier way. The trick is to notice that multiplying both sides of the equation by $\frac{1}{2}$ has the exact same effect as subtracting $\frac{1}{2}$ from both sides:

Multiply both sides by $\frac{1}{2}$:	Subtract $\frac{1}{2}$ from both sides:
$S = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots$	$S = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots$
$\frac{1}{2}S = \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \dots$	$S - \frac{1}{2} = \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \dots$

The right-hand sides of the final line of both equations are the same, so the left-hand sides of both equations must be equal. Taking the left-hand sides of both equations, and solving:

$$\begin{aligned} \frac{1}{2}S &= S - \frac{1}{2} \\ \frac{1}{2}S &= \frac{1}{2} \\ S &= 1 \end{aligned} \quad (1.24)$$

The fact that the infinite series adds up to one tells us that Zeno was wrong. If we keep covering half the distance, but do it an infinite number of times, eventually we will cover the entire distance. The sum of all the half trips equals one full trip.

To generalize Zeno's paradox, assume we have the following series:

$$S = \sum_{i=1}^{\infty} \delta^i \quad (1.25)$$

In Zeno's case, δ was $1/2$. Because the members of the series are all powers of the same constant, we refer to these types of series as geometric series. As long as $|\delta|$ is less than one, the sum will be finite and we can employ the same basic strategy as before, this time multiplying both sides by δ .

$$\begin{aligned} \delta S &= \sum_{i=1}^{\infty} \delta^{i+1} \\ \delta S &= S - \delta^1 = S - \delta \\ S(1 - \delta) &= \delta \\ S &= \frac{\delta}{1 - \delta} \end{aligned} \quad (1.26)$$

Substituting $1/2$ for δ , we see that the general equation agrees with our previously obtained result for Zeno's paradox.

Before deriving Equation 1.26, we stipulated that $|\delta|$ had to be less than one. The reason that $|\delta|$ has to be less than one may not be obvious. If δ is equal to one, we are simply adding together an infinite number of ones, and the sum is infinite. In this case, even though it requires us to divide by zero, Equation 1.26 will produce the correct answer.

If δ is greater than one, the sum is also infinite, but Equation 1.26 will give you the wrong answer. The reason is subtle. If δ is less than one, then δ^∞ converges to zero. When we multiplied both sides of the original equation by δ , in effect we added a $\delta^{\infty+1}$ term to the end of the original equation. If $|\delta|$ is less than one, this term is also zero, and the sum is unaltered. If $|\delta|$ is greater than one, however, this final term is itself infinitely large, and we can no longer assume that the sum is unaltered. If this is at all unclear, wait until the end of the following section on finite series, where we will revisit the issue. If δ is less than -1 , the series will oscillate between increasingly large negative and positive values and will not converge. Finally, if δ equals -1 , the series will flip back and forth between -1 and $+1$, and the sum will oscillate between -1 and 0 .

One note of caution: In certain financial problems, you will come across geometric series that are very similar to Equation 1.25 except the first term is one, not δ . This is equivalent to setting the starting index of the summation to zero ($\delta^0 = 1$). Adding one to our previous result, we obtain the following equation:

$$S = \sum_{i=0}^{\infty} \delta^i = \frac{1}{1 - \delta} \quad (1.27)$$

As you can see, the change from $i = 0$ to $i = 1$ is very subtle, but has a very real impact.

SAMPLE PROBLEM

Question:

A perpetuity is a security that pays a fixed coupon for eternity. Determine the present value of a perpetuity, which pays a \$5 coupon annually. Assume a constant 4% discount rate.

Answer:

$$V = \sum_{i=1}^{\infty} \frac{\$5}{(1.04)^i}$$

$$V = \$5 \sum_{i=1}^{\infty} \left(\frac{1}{1.04} \right)^i = \$5 \sum_{i=1}^{\infty} 0.96^i = \$5 \frac{0.96}{1 - 0.96} = \$5 \cdot 25$$

$$V = \$125$$

Finite Series

In many financial scenarios—including perpetuities and discount models for stocks and real estate—it is often convenient to treat an extremely long series of payments as if it were infinite. In other circumstances we are faced with very long but clearly finite series. In these circumstances the infinite series solution might give us a good approximation, but ultimately we will want a more precise answer.

The basic technique for summing a long but finite geometric series is the same as for an infinite geometric series. The only difference is that the terminal terms no longer converge to zero.

$$\begin{aligned}
 S &= \sum_{i=0}^{n-1} \delta^i \\
 \delta S &= \sum_{i=0}^{n-1} \delta^{i+1} = S - \delta^0 + \delta^n \\
 S &= \frac{1 - \delta^n}{1 - \delta}
 \end{aligned} \tag{1.28}$$

We can see that for $|\delta|$ less than 1, as n approaches infinity δ^n goes to zero and Equation 1.28 converges to Equation 1.26.

In finance, we will mostly be interested in situations where $|\delta|$ is less than one, but Equation 1.28, unlike Equation 1.26, is still valid for values of $|\delta|$ greater than one (check this for yourself). We did not need to rely on the final term converging to zero this time. If δ is greater than one, and we substitute infinity for n , we get:

$$S = \frac{1 - \delta^\infty}{1 - \delta} = \frac{1 - \infty}{1 - \delta} = \frac{-\infty}{1 - \delta} = \infty \tag{1.29}$$

For the last step, we rely on the fact that $(1 - \delta)$ is negative for δ greater than one. As promised in the preceding subsection, for δ greater than one, the sum of the infinite geometric series is indeed infinite.

SAMPLE PROBLEM

Question:

What is the present value of a newly issued 20-year bond, with a notional value of \$100, and a 5% annual coupon? Assume a constant 4% discount rate, and no risk of default.

Answer:

This question utilizes discount factors and finite geometric series.

The bond will pay 20 coupons of \$5, starting in a year's time. In addition, the notional value of the bond will be returned

with the final coupon payment in 20 years. The present value, V , is then:

$$V = \sum_{i=1}^{20} \frac{\$5}{(1.04)^i} + \frac{\$100}{(1.04)^{20}} = \$5 \sum_{i=1}^{20} \frac{1}{(1.04)^i} + \frac{\$100}{(1.04)^{20}}$$

We start by evaluating the summation, using a discount factor of $\delta = 1/1.04 \approx 0.96$:

$$S = \sum_{i=1}^{20} \frac{1}{(1.04)^i} = \sum_{i=1}^{20} \left(\frac{1}{1.04} \right)^i = \sum_{i=1}^{20} \delta^i = \delta + \delta^2 + \dots + \delta^{19} + \delta^{20}$$

$$\delta S = \delta^2 + \delta^3 + \dots + \delta^{20} + \delta^{21} = S - \delta + \delta^{21}$$

$$S(1 - \delta) = \delta - \delta^{21}$$

$$S = \frac{\delta - \delta^{21}}{1 - \delta} = 13.59$$

Inserting this result into the initial equation we obtain our final result:

$$V = \$5 \times 13.59 + \frac{\$100}{(1.04)^{20}} = \$113.59$$

Note that the present value of the bond, \$113.59, is greater than the notional value of the bond, \$100. In general, if there is no risk of default, and the coupon rate on the bond is higher than the discount rate, then the present value of the bond will be greater than the notional value of the bond.

When the price of a bond is less than the notional value of the bond, we say that the bond is selling at a discount. When the price of the bond is greater than the notional, as in this example, we say that it is selling at a premium. When the price is exactly the same as the notional value we say that it is selling at par.