

# Financial Modeling for Managers

with Excel<sup>®</sup> applications

second edition

**Dawn E. Lorimer**  
**Charles R. Rayhorn**

**AUTHORS**  
ACADEMIC PRESS

The logo for Academic Press, featuring the word "AUTHORS" in a large, bold, serif font above the words "ACADEMIC PRESS" in a smaller, bold, sans-serif font. Below the text is a stylized graphic consisting of three horizontal lines that curve upwards on the left side, resembling a book spine or a stylized 'C'.

Copyright 2002 by Authors Academic Press

**Financial Modeling for Managers  
with Excel Applications**

All rights reserved. Previous edition copyright 1999 by Dunmore Press Limited. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

ISBN 0-9703333-1-5

Acquisition Editors: *Trond Randoy, Jon Down, Don Herrmann*

Senior Production Manager: *Cynthia Leonard*

Layout Support: *Rebecca Herschell*

Marketing Manager: *Nada Down*

Copyeditor: *Michelle Abbott*

Cover Design: *Tom Fenske and Cynthia Leonard*

Cover Photo: *London Dealing Room. Courtesy National Australia Bank, London.*

Printer: *EP Imaging Concepts*

Acknowledgments

Robert Miller—Northern Michigan University, Michael G. Erickson—Albertson College, and Jacquelynne McLellan—Frostburg State University all contributed greatly to the production of this book through their thoughtful reviews.

When ordering this title, use ISBN 0-9703333-1-5



[www.AuthorsAP.com](http://www.AuthorsAP.com)

Library of Congress Cataloging-in-Publication Data

Lorimer, Dawn E., 1950-

Financial modeling for managers : with Excel applications / Dawn E.

Lorimer, Charles R. Rayhorn.— 2nd ed.

p. cm.

Rev. ed. of: Financial maths for managers / Dawn E. Lorimer, Charles R.

Rayhorn. c1999.

Includes bibliographical references and index.

ISBN 0-9703333-1-5 (pbk.)

1. Business mathematics. 2. Microsoft Excel for Windows. 3.

Electronic spreadsheets. 4. Financial futures. I. Rayhorn, Charles R.,

1949- II. Lorimer, Dawn E., 1950- Financial maths for managers. III.

Title.

HF5691 .L58 2001

650'.01'513—dc21

2001005810

# CONTENTS

Preface .....	1
<b>Part I. Interest rates and foreign exchange</b>	
<hr/>	
Chapter 1: <i>Compounding and discounting</i> .....	<b>7</b>
1.1 Notation and definitions .....	8
1.2 Time value of money .....	8
1.3 Simple interest, bills and other money market securities .....	10
1.4 Compound interest .....	16
1.5 Linear interpolation .....	19
1.6 Real interest rates .....	20
Chapter 2: <i>The valuation of cash flows</i> .....	<b>31</b>
2.1 Notation and definitions .....	32
2.2 Cash flow representation .....	34
2.3 Valuing annuities .....	36
2.4 Quotation of interest rates .....	39
2.5 Continuous time compounding and discounting .....	44
2.6 Fixed interest securities .....	47
2.7 Duration and value sensitivity .....	57
2.8 Interest rate futures .....	61
Chapter 3: <i>Zero's, forwards, and the term structure</i> .....	<b>73</b>
3.1 Notation and definitions .....	74
3.2 A brief mystery of time .....	75
3.3 Zero coupon rates .....	76
3.4 Implied forward rates .....	79
3.5 Forwards, futures and no-arbitrage .....	81
3.6 Computing zeros and forwards .....	86
3.7 Algorithm: computing zeros and forwards from swap data .....	89
3.8 Concluding remarks .....	97

Chapter 4: <i>FX Spot and forwards</i>	<b>103</b>
4.1 Notation and definitions .....	104
4.2 Spot exchange rate quotations .....	105
4.3 Inversions .....	108
4.4 Cross rates .....	109
4.5 Money market forward rates .....	114

## Part II. Doing it the Excel® way

---

Chapter 5: <i>Learning by doing: an introduction to financial spreadsheets</i>	<b>127</b>
5.1 Step-by-step bond valuation example .....	128
5.2 Do-it-yourself fixed interest workshop .....	137

## Part III. Statistical analysis and probability processes

---

Chapter 6: <i>Statistics without becoming one</i>	<b>151</b>
6.1 Notation and definitions .....	152
6.2 Introduction to data .....	153
6.3 Downloading data .....	154
6.4 Data exploration .....	156
6.5 Summary measures .....	159
6.6 Distribution function and densities .....	164
6.7 Sampling distributions and hypothesis testing .....	168
Chapter 7: <i>Regression and financial modeling</i>	<b>175</b>
7.1 Notation and definitions .....	176
7.2 Bivariate data exploration .....	177
7.3 Regression statistics .....	182
7.4 More regression theory: goodness of fit .....	187
7.5 The CAPM beta .....	189
7.6 Regression extensions: multiple regression .....	196

Chapter 8: *Introduction to stochastic processes* **203**

8.1	Notation and definitions .....	204
8.2	Expectations .....	205
8.3	Hedging .....	206
8.4	Random walks and Ito processes. ....	213
8.5	How Ito processes are used .....	217
8.6	Volatility models .....	223
8.7	Other time series buzzwords .....	226

Part IV. Many variables

---

Chapter 9: *Many variables* **235**

9.1	Vectors and matrices .....	236
9.2	Matrix inverses and equation solving .....	240
9.3	Statistics with matrices .....	243
9.4	Portfolio theory with matrices .....	245
9.5	Practicum .....	248

Part V. Appendix

---

*General Mathematical review* **264**

A1	Order of operations .....	264
A2	Multiplication and division with signed numbers .....	265
A3	Powers and indices .....	267
A4	Logarithms .....	272
A5	Calculus .....	273



# Preface

## WHY READ THIS BOOK?

We wrote the book because we felt that financial professionals needed it; and those who are training to be financial professionals, as students in colleges, will need it by the time they finish. Not all of us are cut out to be Quants, or would even want to. But pretty much all of us in the financial world will sooner or later have to come to grips with two things. The first is basic financial math and models. The second is spreadsheeting. So we thought that the market should have a book that combined both.

The origins of the book lie in financial market experience, where one of us (DEL), at the time running a swaps desk, had the problem of training staff fresh from colleges, even good universities, who arrived in a non operational state. Rather like a kitset that had to be assembled on the job: you know the bits are all there, but they can't begin to work until someone puts all the bits together. Their theoretical knowledge might or might not have been OK, but instead of hitting the ground running, the new graduates collapsed in a heap when faced with the "what do I do here and now?" problem. The present book evolved out of notes and instructional guidelines developed at the time. In later versions, the notes took on an international flavor and in doing so, acquired a co author (CH), to become the present book, suitable for the U.S. and other international markets.

The material you will find here has been selected for maximal relevance for the day to day jobs that you will find in the financial world, especially those concerned with financial markets, or running a corporate treasury. Everything that is here, so far as topics are concerned, you can dig up from some theoretical book or journal article in finance or financial economics. But the first problem is that busy

professionals simply do not have the time to go on library hunts. The second problem is that even if you do manage to find them, the techniques and topics are not implemented in terms of the kind of computational methods in almost universal use these days, namely Excel or similar spreadsheets.

No spreadsheet, no solution, and that is pretty well how things stand throughout the financial world. Students who graduate without solid spreadsheeting skills automatically start out behind the eight ball. Of course, depending on where you work, there can be special purpose packages; treasury systems, funds management systems, database systems, and the like, and some of them are very good. But it is inadvisable to become completely locked into special purpose packages, for a number of reasons. For one thing, they are too specialized, and for another, they are subject to “package capture”, where firms become expensively locked into service agreements or upgrades. So in our own courses, whether at colleges or in the markets, we stress the flexibility and relative independence offered by a multipurpose package such as Excel or Lotus. Most practitioners continue to hedge their special systems around with Excel spreadsheets. Others who might use econometric or similar data crunching packages, continue to use Excel for operations like basic data handling or graphing. So, spreadsheets are the universal data and money handling tool.

It is also a truism that the financial world is going high tech in its methodologies, which can become bewildering to managers faced with the latest buzzwords, usually acronymic and often incomprehensible. This leaves the manager at a moral disadvantage, and can become quite expensive in terms of “consultant capture”. In writing the book we also wanted to address this credibility gap, by showing that some, at least, of the buzzwords can be understood in relatively plain terms, and can even be spreadsheeted in one or two cases. So without trying to belittle the quants, or the consultants who trade in their work, we are striking a blow here for the common manager.

It remains to thank the many people who have helped us in preparing the book. As the project progressed, more and more people from the financial and academic communities became involved, and some must be singled out for special mention. Larry Grannan from the Chicago Mercantile Exchange responded quickly to our questions on CME futures contracts. Cayne Dunnett from the National Bank of Australia (NAB) in London gave generously of his time to advise us on financial calculations. Together with Ken Pipe, he also organized the photo shoot for the front cover, taken in the NAB’s new London dealing room. Joe D’Maio from the New York office of the NAB pitched in with assistance on market conventions and products. Andy Morris at Westpac in New York provided helpful information on US financial products. Penny Ford from the BNZ in Wellington, New Zealand kindly assisted with technical advice and data.



Among the academic community, Jacquelynne McLellan from Frostburg State University in Maryland and Michael G. Erickson from Albertson College in Ohio read the entire manuscript and took the trouble to make detailed comments and recommendations. Roger Bowden at Victoria University of Wellington read through the manuscript making many helpful suggestions, and kept us straight on stochastic processes and econometric buzzwords. Finally, it has been great working with Cynthia Leonard and Tom Fenske of Authors Academic Press. They have been encouraging, patient, and responsive at all times, which has made this project a positive experience for all.



# Part One

## INTEREST RATES AND FOREIGN EXCHANGE

In the world of finance, interest rates affect everyone and everything. Of course, this will be perfectly obvious to you if you work in a bank, a corporate treasury, or if you have a home mortgage. But even if you work in equities management – or are yourself an investor in such – you will need to have some sort of familiarity with the world of fixed interest: the terms, the conventions and the pricing. When interest rates go up, stock prices go down; and bonds are always an alternative to equities, or part of a portfolio that might include both.

So the conventions and computations of interest rates and the pricing of instruments that depend on interest rates are basic facts of life. We have another agenda in putting the discussion of interest rates first. Many readers from the industry will need to get back into the swing of things so far as playing around with symbols and numbers is concerned, and even students might like a refresher. Fixed interest arithmetic is an excellent way to do this, for the math is not all that complicated in itself, and the manipulation skills that you need are easily developed without having to puzzle over each step, or feel intimidated that the concepts are so high-powered that it will need ten tons of ginseng to get through it all. Once you have built up a bit of confidence with the basic skills, then you can think about going long in ginseng for the chapters that follow.

As well as interest rates, we have inserted a chapter on the arithmetic of foreign exchange, incorporating the quotation, pricing and trading of foreign currencies for much the same reasons. These days everyone has to know a bit about the subject, and again, the math is not all that demanding, although in a practical situation you really do have to keep your wits about you.

You will notice that Excel is not explicitly introduced in Part 1. We do this in Part 2, where we can use the material of Part 1 to generate some computable examples. In the meantime, it is important that you can execute the interest rate arithmetic on a hand-calculator. Practically any commercial calculator (apart from the simple accounting ones) will have all the functions that you need, and indeed most can be executed using a very basic classroom scientific type calculator – it just takes a bit longer. The beauty of acquiring a true financial calculator is that in addition to special functions like the internal rate of return, it has several storage locations, useful for holding intermediate results when solving more complex problems. At any rate, keep a calculator nearby as you read on.



---

# One

## Compounding and Discounting

This chapter has two objectives. The first is to review some basic interest rate conventions. Here we address questions like, what is the true rate of interest on a deal? Is it what you see in the ad, what the dealer quoted you over the phone, or something different altogether? And how does this rate compare with what is offered elsewhere?

The interest rate jungle is in some respects like shopping for a used car - you can get some good deals, but also some pretty disastrous ones, where the true cost is hidden beneath a fancy PR package. Unless you know what you are doing, things can get pretty expensive.

The second objective is to introduce you to the market ambience where interest rates are quoted. In this chapter, we are largely concerned with money market instruments, which are a particular sort of fixed interest instrument used by corporate treasurers every day. Even if you are not a corporate treasurer, you are surely your own personal treasurer, and knowing about these instruments will help you in your own financing and investment decisions. Studying these instruments will sharpen your understanding of the interest rate concept.

## 1.1 Notation and definitions

APR	Annual percentage interest.
FV	Future value of money, which is the dollar amount expected to be received in the future. It generally includes amounts of principal and interest earnings.
PV	Present value of money, the value today of cash flows expected in the future discounted at an (or in some situations at more than one) appropriate rate.
$P_t$	The price at time $t$ . Can also stand for the principal outstanding at time $t$ .
$i$	Symbol for the effective interest rate, generally expressed as a decimal. Can vary over time, in which case often indexed as $i_t$ .
I	Total money amount of interest earned.
D	Number of days to maturity.
$Dpy$	Number of days per year. Determined by convention, usually 365 or 360 days.
$n$	Number of years (greater than one). Can be fractional.
$m$	Number of compounding periods in a year
C	The amount of the periodic payment or <i>coupon</i> .
$\pi$	The rate of inflation.

## 1.2 Time value of money

There's a saying that time is money. An amount of money due today is worth more than the same amount due some time in the future. This is because the amount due earlier can be invested and increased with earnings by the later date. Many financial agreements are based on a flow (or flows) of money happening in the future (for instance, housing mortgages, government bonds, hire purchase agreements). To make informed financial decisions, it is crucial to understand the role that time plays in valuing flows of money.

## Interest and interest rates

*Interest* is the income earned from lending or investing capital.

The *rate of interest per period* is the amount of interest earned for the period concerned, per unit of *capital* or *principal* invested at the beginning of the period.

Interest is often quoted as a *percent*.

### Example 1.1

If interest of \$15 is payable at the end of a year in respect of an investment or loan of \$200, then the annual rate of interest is  $15/200 = .075$  expressed as a decimal, or  $100\% \times .075 = 7.5\%$ .

To avoid confusion, the decimal form of the interest rate will be used for calculations, except where a formula explicitly calls for the percentage form.

## Nominal, annual percentage, and effective interest rates

An interest rate is usually expressed nominally (the “nominal rate”) as an annual rate, or *percent per annum* (% p.a.). Interest may be calculated either more or less frequently than annually, on a simple or compound basis, and may be required at the beginning of the loan instead of at the end of the loan (known as “discounting the interest”).

Because of these differences and the potential for misleading consumers, Congress enacted the Consumer Credit Protection Act of 1968. This act launched Truth in Lending disclosures that require creditors to state the cost of borrowing using a common interest rate known as the *annual percentage rate* (APR). If the cost of borrowing includes compounding, another interest rate, known as the *effective annual rate*, should be used.

For example, many credit card companies charge approximately 1.5% a month on average monthly balances. The nominal (quoted) rate would also be the APR in this case and would be calculated as  $1.5\% \times 12 = 18\%$ . The effective annual rate

would actually be higher because of compounding, a subject we will discuss in more detail later. The effective rate would be calculated as  $((1 + .015)^{12} - 1) \times 100\%$  or 19.6%. In general, the effective rate can be calculated by replacing 12 by the number of compounding periods ( $m$ ) during the year. Thus, the equation for the effective rate is:

$$((1 + i)^m - 1) \times 100\% \quad (1)$$

Let's look at a situation where the nominal rate and the APR are not the same. The nominal and APR are always different when interest has to be prepaid or when there are fees associated with getting a particular nominal interest rate (e.g. points). Many loans on accounts receivable require the interest on the loan to be deducted from the loan proceeds. This is known as prepaying the interest. For example, let's assume that the nominal rate is 6.00% for one year and that the interest has to be 'prepaid'. If the loan is for \$1000.00, the loan proceeds in this case would be \$940.00, or  $\$1000.00 - .06 \times (\$1000.00)$ . The APR in this example is  $\$60.00/\$940.00$  or about 6.38%. Without truth in lending, the lending institution could claim that the cost of the loan is actually lower than it is. The effective rate would also be 6.38% because there is only one compounding period.

There will be more on nominal and effective rates later. These are important concepts for anyone involved with financial transactions.

### 1.3 Simple interest, bills and other money market securities

The use of simple interest in financial markets is confined mainly to short term transactions (less than a year), where the absence of compounding is of little importance and where the practice of performing calculations quickly, before modern computing aids became widely available, was necessary.

Simple interest can be misleading if used for valuation of long term transactions. Hence, its application in financial markets is usually limited to the calculation of interest on short-term debt and the pricing of money market securities.

When the interest for any period is charged only on the original principal outstanding, it is called *simple interest*. (In this situation, no interest is earned on interest accumulated in a previous period.) That is:



*Simple interest amount = Original Principal × Interest Rate × Term of interest period*

In symbols:  $I = P_0 \times i \times t$  (2)

where  $t = \frac{D}{dpy}$

Here there is no calculation of interest on interest; hence, the interest amount earned per period is constant.

### Future value

In a simple interest environment:

$$FV = P_0 + I = P_0 + (P_0 \times i \times t)$$

That is,  $FV = P_0 (1 + (i \times t))$  (3)

where  $t = \frac{D}{dpy}$

#### Example 1.2

Calculate the amount of interest earned on a deposit of \$1m for 45 days at an annual interest rate p.a. of 4.75%. What is the future value of this deposit?

$$\text{Interest} = P_0 \times i \times t = \$1,000,000 \times .0475 \times \frac{45}{365} = \$5,856.16$$

$$\begin{aligned} FV &= P_0 + I = P_0 (1 + (i \times t)) = \$1,000,000 (1 + (.0475 \times \frac{45}{365})) \\ &= \$1,005,856.16 \end{aligned}$$

### Present value

Formula (3) may be rearranged by dividing by  $(1 + (i \times t))$  so that:

$$PV = P_0 = \frac{FV}{(1 + (i \times t))} \quad (4)$$

In this case, the original principal,  $P_0$ , is the *present value* and therefore the *price* to be paid for the FV due after  $t$  years (where  $t$  is generally a fraction) calculated at a yield of  $i$ .

#### Example 1.3

Calculate the PV of \$1m payable in 192 days at 4.95% p.a. on 365-dpy basis.

$$PV = \frac{FV}{(1 + (i \times t))} = \frac{\$1,000,000}{(1 + (.0495 \times 192 \div 365))} = \$974,622.43$$

There are two methods used for pricing money market securities in the US: the *bank discount* and the *bond equivalent yield* approach. Examples of money market securities that are priced using the discount method include U.S. Treasury Bills, Commercial Paper, and Bankers' acceptances. Formula (4) can't be used directly for valuing these securities because of the particular rate, the bank discount rate, which is usually quoted (see below).

Examples of money market securities that are discounted using the bond equivalent yield approach include Certificates of Deposit (CD's), repos and reverses, and Federal Funds. Also, short dated coupon-paying securities with only one more coupon (interest payment) from the issuer due to the purchaser, and floating rate notes with interest paid in arrears can fit into this category. For money market instruments using the bond equivalent yield, Formula (4) can be used directly.

There are two key differences involved in these pricing methods. The bond equivalent yield uses a true present value calculation and a 365-day year. It applies an interest rate appropriately represented as the interest amount divided by the starting principal. The bank discount method uses a 360-day year and it does not use a normal present value calculation. The interest rate in this case is taken as the

difference between the FV and the price of the instrument divided by the FV. Examples and solutions are provided in the following sections.

The appendix to this chapter describes the most frequently used money market instruments.

### **Pricing a security using the bond equivalent method**

*Pricing a discount security per \$100 of face value.*

$P$  = purchase price (present value)  
 $FV$  = value due at maturity (also usually the face value of the security)  
 $i$  = interest rate at which the security is purchased  
 $t$  = D  
       *dpy*

*dpy* = 365

$$P = \frac{FV}{(1 + (i \times t))} = \frac{100}{(1 + (i \times t))}$$

For a short dated coupon-paying security with only one more coupon (interest payment) the FV becomes: \$100 face value + \$C coupon payment.

Example 1.3, above, illustrates the bond equivalent method.

### **Pricing a security using the bank discount method**

*Pricing a discount security per \$100 of face value.*

$P$  = purchase price (present value)  
 $FV$  = value due at maturity (also the face value of the security)  
 $i_{BD}$  = bank discount rate =  $\frac{FV - P}{FV} \times \frac{1}{t}$   
 $t$  = D  
       *dpy*

*dpy* = 360

$$P = FV \times (1 - (i_{BD} \times t)) = 100 \times (1 - (i_{BD} \times t)) \quad (5)$$

Securities priced using the bank discount approach use a non-present value equation because the interest rate,  $i_{BD}$ , is not a typical return. It is a gain (FV-P) divided not by the starting point price (P), but by the ending value (FV). Multiplying by 360 days annualizes it.

In order to use the standard present value calculation (the bond equivalent yield approach), the bank discount interest rate,  $i_{BD}$ , must be converted into a normal interest rate (known as the bond equivalent yield). The equation for doing this is

$$i = \frac{365 \times i_{BD}}{360 - (D \times i_{BD})}$$

#### Example 1.4

The discount rate for the March 15, 2001 T-Bill (17 days until maturity) quoted in the February 26, 2001 *Wall Street Journal* was 5.14%. Let's calculate the price (PV) for \$10,000 of face value.

$$PV = \$10,000 \times (1 - (.0514 \times \frac{17}{360})) = \$9975.73$$

You could convert the discount rate, 5.14%, to the bond equivalent yield and use the bond equivalent approach shown in Example 1.3 using Formula (3). The bond equivalent rate for the bank discount rate is 5.224% =

$$\frac{365 \times .0514}{360 - (17 \times .0514)}$$

$$PV = \frac{FV}{(1 + (i \times t))} = \frac{\$10,000}{(1 + (.05224 \times 17 \div 365))} = \$9,975.73$$

(Note: These calculations were done rounding to the 16<sup>th</sup> place. You might get a different result if you round to fewer places. These intermediate calculations should be taken to at least seven decimal places.)

#### Money Market Yields

There is yet another method for pricing short term money market securities of which players in the US market need to be aware. This method, common in the

Eurodollar markets, uses money market yields with interest calculated as:

$$\text{Interest} = \text{Face Value} \times [i_{MMY} \times d / 360]$$

The price of such an instrument is calculated using the same technique as the bond equivalent method, but the days per year ( $dpy$ ) is taken to be 360 days.

### A note on market yield

Short-term securities are quoted at a rate of interest *assuming that the instrument is held to maturity*. If the instrument is sold prior to maturity, it will probably achieve a return that is higher or lower than the yield to maturity as a result of capital gains or losses at the time it is sold. In such cases, a more useful measure of return is the *holding period yield*.

### Holding period yield (for short-term securities)

The holding period yield ( $Y_{hp}$ ) is the yield earned on an investment between the time it is purchased and the time it is sold, where that investment is sold prior to maturity. For short-term securities, this is calculated as:

$$Y_{hp} = \frac{P_{sell} / P_{buy} - 1}{t}, \quad (6)$$

where  $t$  is defined above (a fraction of a year),  $P_{sell}$  is the price at which you sold the security, and  $P_{buy}$  is the price at which you bought it.

#### Example 1.5

Suppose the investor from Example 1.4 sold the T-bill (previously purchased at a bank discount rate of 5.14%) at a new discount rate of 5.10% when the bill had just three days to run to maturity. The selling price is calculated:

$$PV = \$10,000 \times \left(1 - \left(.0510 \times \frac{3}{360}\right)\right) = \$9995.75$$

The holding period yield p.a. is:

$$Y_{hp} = \frac{(9,995.75 / 9975.73) - 1}{(14 / 365)} = 0.0523 = 5.23\%$$

## 1.4 Compound interest

Compound interest rates mean that interest is earned on interest previously paid. Bonds are priced on this basis and most bank loans are as well. Many deposit accounts also have interest calculated on a compound basis. Before we formalize things, let's start with a few examples.

### Example 1.6

\$100 invested at 5% annual interest would be worth \$105 in one year's time.

If the same investment is held for two years at 5% and the interest compounds annually, the future value of the investment is:

$$FV = \$105 \times (1 + .05) = \$110.25$$

$$\text{That is, } FV = \$100 \times (1.05) \times (1.05) = \$100 \times (1.05)^2$$

At the end of three years, assuming the same compound interest rate, the investment would have an accumulated value of:

$$\$100 \times (1.05)^3 = \$115.76$$

In general, the formula for *accumulating* an amount of money for  $n$  periods at effective rate  $i$  per period (or for calculating its *future value*) is:

$$AV = FV = P_0 \times (1 + i)^n \quad (7)$$

$$\text{or } FV = PV \times (1 + i)^n \quad (\text{again, } AV = FV)$$

Note that this formula refers to  $n$  periods at effective rate  $i$  per period. Thus,  $n$  and  $i$  can relate to quarterly, monthly, semi annual, or annual periods.

It is easy to see that Formula (7) can be rearranged to give a formula for the present value at time 0 of a payment of  $FV$  at time  $n$ .

$$PV = P_0 = \frac{FV}{(1+i)^n} = FV \times (1+i)^{-n} \quad (8)$$

The following problems illustrate how to find the *future* or *accumulated value* and the *present value* of single cash flows.

**Example 1.7**

Find accumulated value of \$1000 after 5 1/4 years at 6.2% per annum compound interest.

**Solution 1.7**

$$P_0 = 1000, \quad n = 5.25, \quad i = 0.062$$

$$AV = P_0(1 + i)^n$$

$$\begin{aligned} AV &= 1000(1.062)^{5.25} \\ &= 1000(1.371367) \\ &= \$1,371.37 \end{aligned}$$

**Example 1.8**

Find the present value of \$1,000,000 due in three years and 127 days from today where interest compounds each year at 6.25% p.a. Assume a 365-day year, so that 127 days is equal to 0.3479452 of a year.

**Solution 1.8**

$$\begin{aligned} PV &= \$1,000,000 (1.0625)^{-3.3479452} \\ &= \$816,304.43 \end{aligned}$$

In some cases we may wish to value cash flows to a specific date that is neither the end date of the final cash flow nor the current date. The following example illustrates the technique for calculating an accumulated value or future value to a specific date prior to the final cash flow.

**Example 1.9**

Find the value at one year and 60 days from today of \$1 million due in two years and 132 days from today, where interest compounds at 6.5% p.a. Assume a 365-day year.

**Solution 1.9**

Preliminary calculations:

$$132 / 365 = 0.36164384$$

$$60 / 365 = 0.16438356$$

$$PV = P_0 = FV(1 + i)^{-n} = \$1,000,000 (1.065)^{-2.36164384}$$

$$\begin{aligned} \text{Value at time } t &= P_0(1 + i)^t \\ &= \$1,000,000(1.065)^{-2.36164384}(1.065)^{1.16438356} \\ &= \$1,000,000(1.065)^{-(2.36164384-1.16438356)} \end{aligned}$$

$$\begin{aligned} \text{Hence, the value at time } t &= V_t = P_0(1 + i)^t = AV(1 + i)^{-n} (1 + i)^t \\ &= AV(1 + i)^{-(n-t)} \end{aligned}$$

**Multiple interest rates**

For money market funds, NOW accounts, sweep accounts (a variation on money market funds and NOW accounts), and other investments, interest is payable on the previous balance (interest and principal) at the prevailing market interest rate. Since the interest rate may change from period to period, multiple rates of interest might apply, but we'll stick with basic concepts for now.

**Example 1.10**

Find the future (i.e. accumulated) value of \$10,000 invested at 6% compound for two years, and 7% p.a. compound interest for the following four years.

$$FV = \$10,000(1.06)^2(1.07)^4 = \$14,728.10$$




## 1.5 Linear interpolation

In many financial situations, it is necessary to estimate a particular value that falls between two other known values. The method often used for estimation is called *linear interpolation*. (There are other forms of interpolation; however, they are beyond the scope of this book.)

### Example 1.11

Suppose that you know the interest rates at two maturity points on a yield curve and are trying to estimate a rate that falls at some maturity between these two points.




**Figure 1.1** Linear interpolation

The yield curve here is clearly not a simple linear function, but a linear approximation between two relatively close points will not be too far off the curve.

We know that the one-year rate is 4.2% and the two-year rate is 5.3%. The linear interpolation for 1.5 years will be 4.75%, or  $(4.2 + 5.3)/2$ . Where the desired interpolation is not the mid-point between the two known points, the following weighting can be applied.

Using the above two points, suppose that we wish to find an interpolated rate for one year and 40 days (or 1.109589 years). A picture often clarifies our thinking, so let  $i_{1.109589}$ , the rate we are seeking, be equal to  $\Delta\%$ , and draw a timeline for  $i_t = f(\text{time})$ .



### Figure 1.2 Timeline

$$\text{Length AB} = 1.109589 - 1.0 = .109589$$

$$\text{Length BC} = 2.0 - 1.109589 = .890411$$

$$\text{Length AC} = 2.0 - 1.0 = 1.0$$

$$\text{The ratio } \frac{AB}{AC} = \frac{.109589}{1.0} = .109589$$

$$\text{and } \frac{BC}{AC} = \frac{.890411}{1.0} = .890411$$

The above problem is like the reverse of a seesaw. When you have players of different weights the heavier one must sit closer to the balancing point of the seesaw. Here we have shifted the fulcrum of the seesaw and must place the heavier weight on the shorter end to achieve a balance. The above ratios are treated as the weights that we apply to find  $\Delta$  as follows:

$$\Delta = \{(4.2\% \times .890411) + (5.3\% \times .109589)\} = 4.3205479\%$$

The interpolated rate  $i$  for one year and 40 days that we have called  $\Delta$  is therefore equal to 4.3205% (rounded to four decimals).

The above interpolation method is used for a number of financial problems.

It can be applied, for example, to the selection of hedge ratios where more than one instrument is used to hedge the underlying exposure.

## 1.6 Real interest rates

In Section 1.2 of this chapter, we introduced the concept of interest and interest rates without giving their economic purpose. Everybody knows that a dollar today is not the same as a dollar tomorrow because we can invest this dollar to earn

interest. But why is interest offered? There are many reasons and entire books have been written on the subject. We will briefly describe three.

First, it is assumed that people would rather consume their incomes now rather than in the future. To induce people to save part of their income and forgo current consumption, a monetary inducement must be paid. The more you enjoy current consumption, the more you must be paid to save. The incentive that you are paid to save is the dollar amount of interest, which as we already know can be expressed as an interest rate. This preference for current consumption is known as your time preference for consumption. The higher the time preference, the higher the interest rate needed to encourage savings.

The second dimension to this story is *risk*, which comes in various flavors. One aspect is the likelihood that your savings will be paid back (along with the interest). The more uncertain you are that you will be paid back (the probability of loss, or risk), the higher the monetary inducement you will demand in order to save. It is beyond the scope of this book to talk about the many aspects of risk, but the underlying principle is that the higher the risk, the higher the interest (or monetary inducement) demanded.

The final dimension of why interest rates exist is *inflation*. Inflation is the increase in the level of consumer prices, or a persistent decline in the purchasing power of money. Suppose that you had a zero time preference for consumption, so you didn't demand interest for saving. Further suppose that you knew with certainty that the future cash flow would materialize, meaning again, you don't require interest to save. In such a scenario, you would still have to charge interest equal to the rate of inflation, just to stay even.

Taken together, these three dimensions make up what is known as the *nominal rate of interest*. When we remove one of these dimensions—the last one, inflation—we are talking about the *real rate of interest*. The real rate of interest is the excess interest rate over the inflation rate, which can be thought of as the purchasing power derived from an investment.

Suppose the rate of inflation is 10%. How much in today's (time 0) purchasing power is \$1 worth when received in the future (time 1)? The answer is  $1/1.1 = \$0.91$  (notice that this is Equation (9), where  $n = 1$ ). Now suppose that you earned 10% interest over the unit time period. Your real command over goods and services at the end of the period would be 1.1 times 0.91, or \$1 worth of purchasing power—exactly the same as what you started with. The only way this can happen is if the *real* rate of interest you have earned is zero percent. What you gained on the nominal rate of interest of 10%, you lost on the devalued dollar at the end of the period. In this case, the purchasing power derived from this investment is zero.

Let us generalize this a bit. Let  $\pi$  be the rate of inflation, and let  $i$  be the nominal rate of interest. Then the real rate  $i_*$  is defined as:

$$1 + i_* = \frac{(1+i)}{(1+\pi)}$$

$$i_* = \frac{(1+i)}{(1+\pi)} - 1 \quad (9)$$

In other words, the real compounding factor  $(1+i_*)$  is equal to the nominal compounding factor discounted by the rate of inflation.

This is the technical definition, and the one that ought to be used when you operate in discrete time, as we did here. You sometimes see an alternative definition of the real rate as simply the nominal rate less the rate of inflation. With a bit of rearranging, Equation (9) can be written:

$$i_* = (i - \pi) \left[ 1 - \frac{\pi}{(1+\pi)} \right] \approx i - \pi. \quad (10)$$

The sign  $\approx$  means, “approximately equal to.” We are saying that the real rate is approximately equal to the difference between the nominal rate  $i$ , and the rate of inflation  $\pi$ . You can see that the approximation works only if the rate of inflation is low; otherwise, the second term in the square brackets is not small. This alternative definition is convenient because it can be calculated in your head and gives you a quick idea of the real rate of interest. However, it is only an approximation.

The rate of inflation is just the percentage change of some suitable price index. Choosing which version of the price index to use—the consumer price index or the GDP deflator—is a point of debate. In this case, we are talking about changes in real wealth and need some sense of what that wealth can buy in terms of consumer goods and services, so the consumer price index is probably the better choice. However, the GDP deflator or some other measure of inflation may be more appropriate in a different decision problem.

This may not mean much to you if you are not an economist. It’s probably true that it’s the economists who are most interested in the real rate of interest. However, sometimes capital budgeting problems, personal financial planning problems, and the like, are cast in terms of real future cash flows, meaning that they are adjusted for the reduced purchasing power of money at those future dates. In such cases, the real rate of interest must be used in discounting. The general rule is: *Like discounts like*. If your cash flows are nominal, use nominal rates; if your cash flows are real, use real rates.

A similar principle applies where the cash flows are measured on an after-tax basis. In this case, you have to use an *after-tax rate of interest*, which is usually calculated by multiplying the nominal rate by  $(1 - t)$ , where  $t$  is the applicable rate of tax. Again, this is an approximation, valid if the product  $i \times t$  is negligible in size. You might have fun trying to define a “real after-tax rate of interest,” but don’t ask us!

At any rate, the nominal rate of interest will be used throughout this book and is used in almost all market dealings.

# Money Market Instruments

## Appendix 1

### U.S. Treasury Bills (T-Bills)

T-Bills are direct obligations of the United States federal government. They are issued at a discount for periods of three months (91 days), six months (182 days), or one year (52 weeks). Since Treasury bills are the most marketable of all money market securities they have been, and continue to be, a popular investment for short-term cash surpluses. T-Bills can be purchased at inception via a direct auction or in an active secondary market that provides instant liquidity. The income earned on T-Bills is taxed at the federal level only. They are sold at a discount and the price is calculated using the bank discount method.

### Certificates of Deposit (CD)

A certificate of deposit is a time deposit with a bank, with interest and principal paid to the depositor at the end of the fixed term of the CD. The denominations of CDs are at will, ranging from \$500 to over \$100,000. As with other deposits in banks, the Federal Deposit Insurance Corporation insures CD's up to the first \$100,000. For smaller (retail) CDs of \$100,000 or less, banks can impose interest penalties where depositors wish to withdraw funds prior to maturity of the contract. Large (wholesale) CDs, over \$100,000, cannot be withdrawn on demand but are readily marketable, while smaller CDs are not. The secondary market for these larger CD's thins as the maturity lengthens. The price of marketable CD's is the present value of the principal and interest at maturity using the bond equivalent yield approach.

### Bankers' Acceptances

Bankers acceptances can best be regarded as IOUs that have been guaranteed by a bank. Businesses or individuals raising funds issue these to a set face value promising payment by a certain date and take it along to their bank who 'endorse' it, meaning that they accept repayment responsibility if the business cannot or will not meet its obligation. This endorsement is the acceptance. There is an active secondary market for these because of their safety (low default risk). They are sold at a discount and the price is calculated using the bank discount method.

**Eurodollars**

Eurodollars are dollar-denominated deposits at foreign branches of U.S. banks, or are dollar-denominated deposits at foreign banks anywhere in the world. Most Eurodollar accounts are large time deposits and are not marketable. A variation of the U.S. CD is the Eurodollar CD, which is marketable and widely quoted in one-month, three-month, and six-month maturities. As with their U.S. counterparts, the price of Eurodollar CD's is the present value of the principal and interest at maturity using the bond equivalent yield approach. Eurodollars are quoted on a money market yield basis with price calculated using a *360-dpy* basis so that:

$$P = \frac{FV}{(1 + i \times t)} = \frac{FV}{\left(1 + i \times \frac{\text{no. of days}}{360}\right)}$$

**Repos and Reverses**

Repurchase agreements, or repos, are usually very short-term – typically overnight – borrowing by securities dealers. Reverse repurchase agreements, or reverses, are overnight lending by securities dealers (i.e. the reverse side of repurchase agreements). They are collateralised by U.S. Treasury securities and are considered very safe. These loans are priced as a package where the value is the present value of the underlying securities using the bond equivalent approach. The difference between the sell price and the buy price will determine the interest rate (using the holding period yield formula). In a repurchase agreement, the lender will transfer same-day funds to the borrower, and the borrower will transfer the Treasury security to the lender – all subject to the provision that the transactions will be reversed at the end of the repo term. Such transactions enable traders to leverage their security holdings.

**Commercial Paper**

Commercial paper is unsecured short-term debt issued by large well-known companies for a period not exceeding 270 days. Maturities longer than 270 days must be registered with the Securities and Exchange Commission and are very rare. The denominations are in multiples of \$100,000 face value and are sold at a discount from face value using the bank discount approach. Commercial paper is sold directly by the issuing corporation or by securities dealers.

## **Fed Funds**

Federal Funds, or Fed Funds, include cash in a member bank's vault and deposits by member banks in the Federal Reserve System. The amount of Federal Funds required is determined by the Federal Reserve and is a percentage of the member banks total deposits. During the normal course of conducting business, some banks at the end of the business day will find themselves below the required percentage, while other banks will find themselves above the required percentage. The bank that is below the required percentage must get up to this percentage. The Fed Funds market exists so that banks in surplus can loan to banks that have a deficit. The loans are usually overnight. This market is only for member banks, but the interest rate that prevails in this market is believed to be the base rate for all short-term rates in the United States and, by extension, much of the world. This rate is one of the key monetary variables that the Federal Reserve Board targets when setting monetary policy.



1. Compute  $P$  (possibly using mathematical Appendix A.1),

$$\text{where } P = (1.06) \times \frac{[1 - ((1.06)^{-2})^3]}{1 - (1.06)^{-2}} .$$

2. Find the annual compound rate at which \$800 will double after eight years.
3. It is known that \$4000 will accumulate to \$4500 after five quarters. Find the quarterly rate of compound interest.
4. With interest at 6% per annum compound, how long will it take \$1000 to accumulate to \$1500?
5. A target retirement benefit of \$200,000 is to be provided after seven years. This benefit will be funded by an initial contribution of \$70,000 and a lump sum payment after five years. If the fund earns 8% per annum compound, determine the lump sum required after five years.
6. A certain vintage car is expected to double in value in 12 years. Find the annual rate of appreciation for this vehicle.
7. How long does it take to triple your money at 10% per annum compound interest? (Hint: Use logs. See mathematical Appendix A.4.)
8. A special “bond” pays \$60,000 in three years and a further \$80,000 in five years from now. Determine the price of the bond assuming that the interest rate for the first three years is 3.5% per half-year compound and 4% per half-year compound thereafter.
9. Calculate the holding period yield on a 90-day Certificate of Deposit (CD) bought at a yield of 5.0% p.a. and sold five days later at 4.8% p.a.
10. Suppose that on April 16, 2001, the following yields prevail in the market for US Treasury notes:

# 1 Exercises

Maturity date	Yield % p.a.
Feb 15, 2004	4.14%
May 15, 2005	4.34%
Aug 15, 2006	4.66%

Using linear interpolation, find the approximate four-year and five-year points on the above yield curve (as of April 16, 2001).

11. If you are familiar with Excel you might like to try the following problem; otherwise you may prefer to attempt answering the question after you have read Chapter 5.

In Question 1 you were asked to compute the value of P. Now suppose you had entered the equation for P into your Excel spreadsheet as:

$$P = (1.06 * (1 - (1.06)^{-2})^3 / 1 - (1.06)^{-2})$$

Would Excel object, and why? How would you fix it?





---

# Two

## The Valuation of Cash Flows

In the last chapter, we saw that T-Bills and other similar financial instruments are single-payment securities. Many financial instruments are not single payments but are a series of extended payments, so we need to modify our valuation concept to include multiple payments. Coupon bonds are a good example of such a series, typically paying a cash flow (coupon) every six months and at the end, the return of the original capital. Thus part of this and other chapters will look at the trading conventions and pricing for bonds. We shall also discover that bonds come in all varieties, and that they can be priced in terms of competing versions of what constitutes an interest rate.

The second part of the chapter shifts the perspective somewhat. Valuing one bond is straightforward enough, once you have decided on the right way to do it. But what about valuing a portfolio of bonds, or controlling the risk of a portfolio against changes in general interest rates? We look at some of the basic tools for doing this, and in particular, the concept of duration, a sensitivity measure used by fixed-interest analysts and managers.

In this chapter we will also consider some of the instruments used in the management of interest rate risk, including T-Bill, Eurodollar and Bond futures contracts. Derivatives of this kind can also be used for open position taking, i.e. speculation. We will look at the terminology, market trading conventions, and pricing.

## 2.1 Notation and definitions

Annuity	A series of equal cash flows occurring at equal time intervals. There are two types of annuities. An <i>ordinary annuity</i> (or deferred annuity) pays cash flows at the end of the period (sometimes referred to as paying in arrears). An <i>annuity due</i> pays cash flows at the beginning of the period.
a	In fractional discounting, number of days from settlement date to the next interest date.
b	Number of days from last interest date to next interest date (i.e. total number of days in the current interest period).
c	Annual coupon rate, expressed as a decimal.
C	The coupon payment ( $FV \times c/2$ for a bi-annual coupon).
D	Macaulay duration: a weighted average of the times to receipt of a specified group of cash flows, the weight being proportional to their present values.
Derivative	An instrument whose price is tied in some way or derived from, the price or yield of another instrument (the underlying 'physical').
EAR	Effective annual interest rate. The true annual rate where a given nominal rate is compounded more than once inside a year.
FV	Nominal or face value of the bond.
$i_t$	The effective interest rate for the period ended at time $t$ , expressed as a decimal. The subscript, $t$ , is omitted if the interest rate is constant each period. In this case, $i$ is the effective rate per period. (As we will see later, this will be equal to $r_m/m$ .)
$r_m$	The nominal annual interest rate (APR) compounded $m$ periods during the year.
$r_1$	Effective rate per annum. The interest rate that would produce the same ending value, as would have resulted if annual compounding had been used in the first place.

n	The number of compounding periods. In the case of an annuity, the number of cash flows. In the case of a coupon bond, the number of coupons remaining until maturity.
P	Market value or price of the bond.
$P_Q$	“Quoted” price, as distinct from invoice price, of a bond.
$PV_i$	The present value of the $i^{\text{th}}$ payment.
$PVIF_{i,n}$	Present value interest factor. This is equal to $\frac{1}{(1+i)^n}$ or $(1+i)^{-n}$ .
$PVIFA_{i,n}$	Present value interest factor for an annuity. This is equal to $\frac{1 - \frac{1}{(1+i)^n}}{i}$ .
PVA	The present value of an ordinary annuity. The letter “A” is added to distinguish the present value of an annuity from other present values. $PVA = C \times PVIFA_{i,n}$ .
$PVA_D$	The present value of an annuity due. The subscript “D” is added to distinguish the present value of an annuity due from ordinary annuities. $PVA_D = C \times PVIFA_{i,n} \times (1+i)$ .
$t_i$	The time at which the $i^{\text{th}}$ payment is made or received.
IRR	Internal rate of return, or the yield to maturity. This is the rate of discount that equates the price (cost) of a bond or bill to the present value of the cash flows—the coupons and the face value payment.
Forward	A forward contract is an agreement calling for future delivery of an asset at a price agreed upon at the inception of the contract.
Futures	A futures contract is a forward contract where the asset is in standardized units, quality and quantity, and which is marked to market (valued) every day until maturity. Futures contracts are traded on exchanges while forward contracts are over the counter instruments – that is, not exchange-traded.

#### Invoice versus quoted price

For a bond, in some markets including the US, the invoice price is

what you actually pay. The quoted price is lower by the amount of interest accrued since the last coupon date.

## 2.2 Cash flow representation


Almost all problems in finance involve the identification and valuation of cash flows. Transactions involving many cash flows can rapidly become complicated. To avoid making errors, it is helpful to follow a systematic method, such as using a cash flow schedule or a timeline diagram to clarify dates, amounts, and sign (+ or -) of each cash flow.

### Example 2.1a

Let's value the following cash flows: \$2.8m is to be paid out at the beginning of year three (this is the same as saying \$2.8m at the end of year 2), \$2.1m is to be received at the end of year four, and \$1.8m is to be received at the end of year five. Calculate the value of these payments today, assuming a compound interest rate of 6.5% p.a. throughout.

#### Timeline diagram

The horizontal line in Figure 2.1 represents time. A vertical line with a positive value cash flow represents an inflow, and a vertical line with a bracketed amount represents an outflow. The timeline and cash flow schedule below illustrate the three cash flows.



**Figure 2.1** Timeline



### Cash flow schedule

Cash flow Time (Years)	Cash flow (\$m)
0	0.00
1	0.00
2	-2.80
3	0.00
4	2.10
5	1.80

At the left is a schedule of the timing and amounts of the cash flows. Inflows are expressed as positive numbers and outflows are expressed as negative numbers. Additional columns could be used to identify the sources and destinations of cash flows. This layout is similar to the way you would set up the problem in a spreadsheet.

Knowing the size, sign and timing of the cash flows, we can proceed with valuing them. Formula (8) from Chapter 1 will be used to solve for the present value of each cash flow. The present values for the three cash flows will then be added together and this sum constitutes the present value of the set of cash flows.

$$PV = \$ -2.8(1.065)^{-2} + 2.1(1.065)^{-4} + 1.8(1.065)^{-5} = \$ 477,518$$

By definition, the present value of a set of cash flows is simply the sum of the present values of the individual cash flows. When everything is valued to the same point in time, you are comparing apples with apples. At times, it may be necessary to value cash flows to a point in time other than the present, that is, to a point in time other than  $t = 0$ . The following example illustrates the valuation of a set of cash flows to a future point in time.

#### Example 2.1b

Calculate the accumulated value of the cash flows in Example 2.1a at time  $t = 3.5$  years (that is, halfway through the fourth year) again, assuming a compound interest rate of 6.5% p.a.

The value halfway through the fourth year is equal to:

$$\text{Value}_{3.5} = PV \times (1.065)^{3.5} = 477,518 \times (1.065)^{3.5} = \$595,269.16$$

## 2.3 Valuing annuities

Many situations arise in practice where a series of equal cash flows occur at regular equal time intervals. Such a series of payments is called an *annuity*. Annuities are very common. Examples include lease payments, insurance premiums, loan repayments, and fixed interest coupons. When a set of cash flows does form an annuity, calculating the present value is relatively easy. As we shall see, it is no longer necessary to value each flow separately.

### Example 2.2

#### The present value of an ordinary annuity

In an *ordinary annuity*, the cash flows fall at the end of each period. Consider a series of \$1 payments at the end of each year for five years at 6% p.a. compounded. Using the definition for the present value of a set of cash flows, we would do the following calculations:

$$\begin{aligned} \text{PVA} &= \$1(1.06)^{-1} + \$1(1.06)^{-2} + \$1(1.06)^{-3} + \$1(1.06)^{-4} + \$1(1.06)^{-5} \\ &= \$1(0.9434 + 0.8900 + 0.83962 + 0.79209 + 0.74726) \\ &= \$ 4.21236 \end{aligned}$$

Thus, for an investment (loan) of \$ 4.21, at 6.0% you would receive (pay) \$1 per year for five years at the end of the period (in arrears).

Notice that for an ordinary annuity, the cash flows are such that when we calculate the present value, all cash flows are brought back (discounted) to one period before the first cash flow is received (or paid out). When we discuss the present value of an annuity due, this will not be the case and will necessitate an amendment to the annuity formula.

#### Ordinary Annuity Formula

To illustrate how the formula is derived, consider Example 2.2 above, except that we shall remain a bit more general with an arbitrary interest rate  $i$ .

$$PVA = \$1[(1+i)^{-1} + (1+i)^{-2} + (1+i)^{-3} + (1+i)^{-4} + (1+i)^{-5}].$$

Multiply both sides of the equation by  $(1+i)$ , and we get:

$$PVA(1+i) = \$1[1 + (1+i)^{-1} + (1+i)^{-2} + (1+i)^{-3} + (1+i)^{-4}].$$

Subtracting the first equation from the second,  $PVA(1+i) - PVA$ :

$$PVA(1+i) - PVA = \$1 \times [1 - (1+i)^{-5}]$$

That is:

$$i \times PVA = \$1 \times \left[ 1 - \frac{1}{(1+i)^5} \right]$$

Then the present value formula for a five-period ordinary annuity where the cash flow is \$1 is given by:

$$PVA = \$1 \times \left[ \frac{1 - \frac{1}{(1+i)^5}}{i} \right]$$

Using the same methodology, we can solve for the present value of any  $n$ -period ordinary annuity with cash flow  $C$  and interest rate  $i$  using the following formula:

$$PVA = C \times \left[ \frac{1 - \frac{1}{(1+i)^n}}{i} \right] \quad (1)$$

The term in brackets is written as  $PVIFA_{i,n}$ , the present value interest factor for an annuity.

**Example 2.3**

How much money would you be able to borrow if you had to repay the loan in 20 years with equal annual payments at an interest rate of 10.5%, if the most you could afford in annual payments is \$3,000? Assume the first payment is in one year's time.

$$PVA = \$3,000 \times \left[ \frac{1 - \frac{1}{(1 + 0.105)^{20}}}{0.105} \right]$$

$$PVA = \$3,000 \times [8.2309089] = \$24,629.73$$

**Present value of an annuity due**

In an *annuity due*, payments are made in advance rather than in arrears. Let us revisit Example 2.2, but move the cash flows to the beginning of the period (or in advance).

**Example 2.4**

Using the definition for the present value of a set of cash flows, we would do the following calculations:

$$\begin{aligned} PVA_D &= \$1(1.06)^0 + \$1(1.06)^{-1} + \$1(1.06)^{-2} + \$1(1.06)^{-3} + \$1(1.06)^{-4} \\ &= \$1(1 + 0.9434 + 0.8900 + 0.83962 + 0.79209) \\ &= \$4.46511 \end{aligned}$$

Thus for an investment (loan) of \$ 4.47, you would receive (pay) immediately \$1 per year for five years, at 6%. Notice that this amount is 1.06 times larger than the \$4.21 we got when we solved this problem before. This is no coincidence. If you compare the above calculation with the one for the ordinary annuity, you will notice that there is one less discounting period for every cash flow. We could take the first set of calculations (for the ordinary annuity) and turn it into the annuity due calculations by multiplying each cash flow by  $(1 + i)$ .

$$(1 + i) PVA = PVA_D$$

You will notice that for an annuity due, the cash flows are such that when we calculate the present value all cash flows are brought back (discounted) to when the first cash flow is received (or paid out). Remember that for the present value of an ordinary annuity, this is not the case.

### Annuity due formula

As indicated above, we could replicate the argument used to derive Equation (1), but the only difference would be that all of the terms would be multiplied by  $(1+i)$ . The resulting formula is:

$$PVA_D = C \times \left[ \frac{1 - \frac{1}{(1+i)^n}}{i} \right] \times (1+i) \quad (2)$$

Be careful when working with annuity problems. The time period to which the cash flows are discounted determines whether you solve the problem with Formula (1) or (2). If all of the cash flows are discounted to one period before the first cash flow, use Formula (1). If all of the cash flows are discounted to when the first cash flow occurs, use Formula (2).

## 2.4 Quotation of interest rates

Interest rates tend to be quoted in many different ways, which sometimes makes comparison between quotes a problem. As we saw in Chapter 1 this was one of the reasons Congress passed the Consumer Credit Protection Act of 1968. Even though we discussed this topic briefly in Chapter 1, a little more detail is necessary.

Compound interest is calculated at the end of each interest period and compounded into (added to) the principal. The number of times interest is compounded in one year is called the *compounding frequency per annum*. The period of time between successive compounding is called the *interest period*.

### Nominal (APR) and effective rates of interest

If \$10,000 is invested at a nominal rate of 12% per annum with monthly compounding, the effective monthly return is 1% compound and the effective annual

return is significantly *higher* than 12%. For at the end of 12 months, the \$10,000 would have grown at 1% compound per month to  $\$10,000(1.01)^{12} = \$11,268.25$ . If you look at the interest accumulation (by subtracting the principal of \$10,000), you can see there is \$1,268.25 of interest earned on a principal of \$10,000. This represents an annual effective rate of 12.6825%. In other words, if you used a 12.6825% p.a. interest rate compounded once during the year you would earn the same as with the 12% compounded monthly in the above example.

This is the meaning of *effective annual interest rate*, EAR. We can see from Table 2.1 below that as compounding frequency is increased, keeping the nominal rate constant, the effective rate increases.

**Table 2.1 Compounding frequency and effective rates**

Compounding Frequency	Nominal Rate % p.a. (APR)	Effective Annual Rate %
Annually	6.00	6.00
Semi-annually	6.00	6.09
Quarterly	6.00	6.13636
Monthly	6.00	6.16778
Daily	6.00	6.18315
Continuous	6.00	6.18365

Up to this point, we have used the letter  $i$  in equations to denote the relevant periodic interest rate. We have used such interest rates in a relatively general way, but now we are going to be more specific and introduce subscripts, so it becomes much clearer to switch to the letter  $r$ .

Consider \$1 invested for a year at the nominal rate of  $r_m$ . This means that the interest compounds  $m$  times a year, and that the effective rate per compounding period when there are  $m$  periods per year is  $r_m/m$ .

By the end of the year, the investment will accumulate to  $\$1(1 + r_m/m)^m$ . This represents the original \$1 invested plus interest of  $(1 + r_m/m)^m - 1$ . Thus, the nominal rate of  $r_m$  is equivalent to an annual effective rate of  $(1 + r_m/m)^m - 1$ .

Consistent with this notation, define  $r_1$  as the effective rate from only one compounding during the year (in other words, just the effective annual rate). We see that  $r_1$  and  $r_m$  are related by the following formula:

$$(1 + r_1) = (1 + r_m/m)^m \quad (3)$$

so that:

$$r_1 = (1 + r_m/m)^m - 1 \quad (4)$$

Let's take a moment and look at Equations (3) and (4). The left hand side of Equation (3) is a one year future value interest factor using the effective annual rate; the right hand side is the one year future value interest factor using the nominal rate with an adjustment for the number of compounding (conversion) periods.

Notice also that Equation (4) above gives the same result as Equation (1) from Chapter 1 because  $r_m/m$  is equal to  $i$ , the periodic rate. Thus, the introduction of  $r_m/m$  formalizes the way in which we calculate the periodic rate introduced in Chapter 1.

### Example 2.5

Find the effective annual rate of interest corresponding to 6% (nominal) p.a. compounded (convertible) quarterly.

Here,  $r_m = 6\%$  p.a. , and  $m = 4$

$$\begin{aligned} \text{Hence from Formula (4), } r_1 &= (1 + .06/4)^4 - 1 \\ &= 1.0613636 - 1 = 6.136\% \end{aligned}$$

As an exercise, you can now check the entries in Table 2.1 above, except for continuous compounding which we will discuss in Section 2.5.

Sometimes it is necessary to work backwards from the effective annual rate of return to find the nominal rate. The next example will show us how to accomplish this.

**Example 2.6**

What nominal rate compounded every month will equate with an effective annual rate of 6.16778%? In other words, solve for  $r_m$ , where  $m = 12$ .

From expression (3) above:  $(1 + r_1) = (1 + r_m/m)^m$

Hence  $(1 + r_1)^{1/m} = (1 + r_m/m)$

$$(1 + r_1)^{1/m} - 1 = r_m/m$$

$$m((1 + r_1)^{1/m} - 1) = r_m$$

So that,  $r_m = m((1 + r_1)^{1/m} - 1)$

In our case:  $r_{12} = 12 ((1 + .0616778)^{1/12} - 1)$

$$= 6\% \text{ p.a. nominal compounding monthly}$$

Of course you probably already knew the answer. Look at Table 2.1 again. This is just the monthly compounding number.

It is market practice to express interest rates as nominal rates per annum. However the nominal rate is strictly meaningless, unless the compounding frequency is specified, or understood. Likewise, specification of the interest period is essential when speaking of effective interest rates.

**Equivalent rates**

Now we introduce a way to convert from one nominal rate, with a specified number of compounding periods in a year, to another nominal rate, with a different number of compounding periods, when they have the same *equivalent annual effective rate*. We will differentiate between two nominal rates by using  $r_m$  to represent one and  $r_k$  to represent another.

Using (3) again:  $(1 + r_1) = (1 + r_m/m)^m$ , and  $(1 + r_1) = (1 + r_k/k)^k$

so  $(1 + r_m/m)^m = (1 + r_k/k)^k$

from which  $r_m = m((1 + r_k/k)^{k/m} - 1)$  (5)



Applying this formula to different nominal rates that do not have the same equivalent annual effective rate will give you nonsense.

### Example 2.7

From Formula (4), we know that the effective rate for 5% nominal p.a. compounded monthly is  $5.11619\% = r_1 = ((1 + .05/12)^{12} - 1) \times 100\%$ . What is the nominal rate compounded quarterly that has the same effective rate?

Using expression (5), we get:

$$r_4 = 4((1 + .05/12)^{12/4} - 1) \times 100\% = 5.02086\%$$

The answer is 5.02086% nominal compounded quarterly. To check this result, try using Formula (4) and compute the effective rate. If the effective rate is the same as above, you have the correct answer.

$$r_1 = (1 + .050286/4)^4 - 1 = .0511619 \text{ or } 5.11619\%$$

The effective rates are the same, so 5% nominal p.a. compounded monthly is equal to 5.02086% nominal p.a. compounded quarterly.

#### Some guidelines on cash flow valuation and effective return:

1. *The effective return must match the cash flow frequency*

When valuing cash flows, remember to match the effective rate per period with the frequency of the cash flows. So, for example, if cash is rolling over monthly, an effective monthly rate should be used for valuation purposes.

2. *When the investment or borrowing period is less than one year, adjust the effective return to an annual basis for comparison purposes.*

**Example 2.8**

Suppose that interest is paid at a nominal rate of 8% p.a. every 90 days. This investment rolls over so that interest is earned on the interest. Assume that it compounds at the same rate for three periods of 90 days each. What is the effective annual return for the investment?

$$\begin{aligned} \text{The effective 90 day rate is} &= 8\% \times 90 / 365 \\ &= 1.9726027\% \\ &= .019726027 \text{ expressed as a decimal} \end{aligned}$$

Compounded for three periods of 90 days at the same rate, the effective rate becomes:

$$\begin{aligned} &= (1.019726027)^3 \\ &= 1.0603531 \end{aligned}$$

That is, the effective return over the 270 days is:

$$= (1.0603531 - 1) \times 100 = 6.0353\%$$

As the investment does not benefit from another compounding period, to find the effective annual rate, simply perform a pro-rata adjustment to 365 days.

$$\text{Effective annual rate} = 6.03531 \times 365 / 270 = 8.1588457\%$$

**2.5 Continuous time compounding and discounting**

Go back to the discussion in Section 2.4 and look at Table 2.1 again. You will notice that as the compounding interval gets smaller, the effective annual rate tends to a limit, which is provided by continuous compounding. This means that there are an infinite number of infinitesimally small compounding periods within the year.

You will recall from Formula (3) or (4) that if there are  $m$  compounding intervals inside a year, then the annual effective compounding factor is

$$(1 + r/m)^m,$$

where we have written the nominal annual rate as  $r$ . To simplify the discussion, we have dropped the subscript  $m$ , which is implied in the rate,  $r$ . In fact, in what follows we are going to keep the rate  $r$  constant as the compounding period  $m$  varies, just as we did in Table 2.1. Of course, the basic unit of time here could be anything we please, but the usual practice is to normalize the time unit as one year. We already know that the effective interest rate is higher than the nominal annual rate  $r$ , and how much greater depends upon the compounding frequency ( $m$ ).

Now, what happens as the compounding intervals diminish: down to a month, then a week, then a day, an hour and so forth? Or equivalently, the number of intervals  $m$  grows larger and larger? There is a mathematical result that says:

$$\lim_{m \rightarrow \infty} \left(1 + \frac{r}{m}\right)^m = e^r.$$

Here the number  $e$  is the base for natural logarithms, one of the most important constants in mathematics. It is approximately equal to 2.71828. Most calculators and all financial calculators have an  $e^x$  key which raises  $e$  to the  $x^{\text{th}}$  power, and also can calculate its reverse: if  $y = e^x$ , then  $x = \ln(y)$ , where the symbol  $\ln$  denotes the natural logarithm.

As you reduce the length of the time compounding interval, you are approaching continuous time where you have an infinite number of infinitesimally small time units. So the math is telling us that the continuous time compounding factor is equal to  $e^r$ , where  $r$  is the nominal annual rate, as  $m$  tends to infinity.

For the nominal annual rate of 12%, you would raise  $e$  to the 0.12 power,  $e^{0.12}$ , and get 1.127497. Subtracting one (1) from this number as in Equation (4) and converting to a percentage, you get the effective annual rate as 12.7497%, the last number in Table 2.1.

What about future values? Recall Equation (7) of Chapter 1, which says that:

$$\text{FV} = P_0 \times (1 + i)^t,$$

where  $t$  is the number of complete years.

If we have more than one compounding period in a year, substitute  $r_m$  for  $i$  to denote the nominal rate. The formula becomes:

$$FV = P_0 \times (1 + r_m/m)^{tm} \quad (6)$$

Equation (6) is the lump sum future value formula when there are  $m$  compounding periods in a year. Letting the units of time become infinitesimally small, and  $m$  infinitely large, changes Equation (6) to:

$$FV = P_0 \times e^{rt} \quad (7)$$

Similarly, the present value formula (Equation (8) in Chapter 1) becomes:

$$PV = P_0 = FV \times e^{-rt} \quad (8)$$

Hence, the exponential function assumes the role of the continuous time compounding or discounting factor. The astute reader will notice that since Chapter 1, we have changed  $n$  to  $t$ . The change is necessary because of the convention of using  $t$  in continuous time problems.

Equations (7) and (8) are functions of  $t$ . Equation (7) has another useful interpretation. To develop it, we shall need a bit of differential calculus (if your math does not extend this far, skip the derivation and just note the final result).

Differentiating Equation (7) and using  $\frac{d}{dt}[e^{rt}] = re^{rt}$ , we get:

$$\begin{aligned} \frac{dFV}{dt} &= rP_0e^{rt} \\ &= rFV, \text{ using (7) again.} \end{aligned}$$

This can be written:

$$\frac{dFV}{dt} \bigg/ FV = r$$

This equation illustrates that the instantaneous proportional rate of change of one's capital is equal to the nominal rate,  $r$ . Or you could say the proportional change is equal to the nominal rate times the small time interval  $dt$ . Our capital is growing at an instantaneous rate of  $r\%$ .

Much of advanced theory in finance is cast in terms of continuous time. There is a saying among mathematicians: "God made the integers, but man made the continuous." If so, man certainly made things easier for himself, because the

mathematics of continuous time turn out to be much more powerful than that for discrete time. We will revisit this later when we discuss stochastic processes.

In the meantime, we note a more practical point. Once you get down to discrete compounding intervals like a day, it turns out that the difference between this and the continuous time valuation is so small that people often choose the ease of continuous time-based calculations.

## 2.6 Fixed interest securities

### Coupon bonds

Coupon bonds are financial instruments that require the issuer to pay:

- A fixed amount, the *face value* (or *principal*), on the maturity or redemption date
- A number of periodic interest payments, known as *coupons*, of an amount that is fixed at the time the bond is issued. It is usual for the coupon to be paid semi-annually.

The market value of a bond should equal the present value of its future cash flows (where the cash flows consist of its face value and the periodic coupons) using the current market yield.

For the discussion that follows, we will return to our original notation, using  $i$  to denote the effective interest rate per period.

The price of a coupon bond ( $P$ ) is given by:

$$P = \frac{C}{(1+i)} + \frac{C}{(1+i)^2} + \dots + \frac{(C + FV)}{(1+i)^n}, \quad (9)$$

where  $C$  stands for the coupon amounts and  $FV$  is the face value of the bond.

When calculating the price of a bond, it is assumed that the bond will be held to maturity, and that *the coupons will be reinvested at the current market yield to maturity*.

This assumption implies that we can rewrite Equation (9) as:

$$P(1+i)^n = C(1+i)^{n-1} + C(1+i)^{n-2} + \dots + C + FV,$$

meaning that a buyer of the bond should be able to achieve an accumulated sum of  $P(1+i)^n$ . In other words, you should be indifferent between accumulating the cost  $P$ , on the one hand, and the stream of rewards, on the other, all at the same rate of interest, namely  $i$ .

For an investor, the above assumption results in a reinvestment risk. If interest rates fall and remain below the purchase market yield, the coupons will be reinvested at a lower rate. Should the investor hold the bond to maturity under such an interest rate scenario, the actual yield achieved by the investor will not match the yield implicit in the price at which the bond was purchased. If interest rates rise and remain above the purchase market yield, the coupons will be reinvested at a higher rate and the actual yield will be higher.

In the United States the market tends to trade bonds on the basis of price. Given a market price  $P$ , Equation (9) determines the effective rate per period ( $i$ ) as the *internal rate of return (IRR)* on the bond. This *IRR* is the per period rate which equates the present value of the cash flows from the bond, including its repayment of principal, with the market price,  $P$ . If the coupons are paid every six months the *IRR* will be a six-month rate and will need to be multiplied by two to annualize it.

In some other parts of the world, such as Australia and New Zealand, bonds tend to be bought and sold on the basis of market yield, which is quoted as a nominal annual rate. As the effective periodic rate,  $i$ , used in Equation (9) is usually semi-annual, you will need to take the quoted nominal annual rate and divide by two to calculate price.

Readers familiar with corporate finance will know that the internal rate of return of a project is the discount rate that equates the present value of the project cash flows with the upfront cost of the project. This is exactly the idea behind Equation (9), where the coupons and final principal payment correspond to the cash flows and the given price of the bond to the upfront cost. In some circumstances involving cash flows of alternating sign, there can be multiple different rates of return from a given set of cash flows, but this awkward property does not feature in bond pricing, where the rate is always unique. As mentioned above, the assumption in using the internal rate of return is that it implicitly assumes that all coupons received are reinvested back into a financial instrument with exactly the same rate of return as the bond. There is no reason this should automatically be true, and in most cases it isn't. Therefore, it can be hazardous to compute the yield from one bond and then use this yield to find the price for another. We'll discuss this further in the next chapter.

The above bond pricing Formula (9) is simple enough to apply if the number of coupon periods left to run is few, or if you have some time and a computer. However, using what we have learned about annuities (Section 2.2), we can take a shortcut in our calculations to find the present value of a coupon bond. The coupon payments valued as an ordinary annuity, along with the present value of the lump sum repayment of principal or face value, determine the total value of the bond. Thus, we can combine Equation (1) from this chapter and Equation (8) from Chapter 1, to get a closed form alternative.

$$P = C \times PVIFA_{i,n} + FV \times PVIF_{i,n} \quad (10)$$

### Example 2.9

Using Formula (10) above, we can price a bond that has four years (eight six-month periods) until maturity. Suppose the bond has a face value of \$100 and pays semi-annual coupons of 12% p.a., and that the market yield for the bond is currently 12% p.a. (compounding semi-annually, i.e., 6% effective per six-months).

The semiannual coupon payments are:  $\$100 \times .12/2 = \$6$

Hence, on the basis of a notional \$100 face value, the price of the bond will be:

$$\begin{aligned} P &= 6 \times PVIFA_{.06,8} + \frac{100}{(1+.06)^8} \\ &= 6 \times \frac{1 - \left(\frac{1}{(1.06)^8}\right)}{0.06} + \frac{100}{(1.06)^8} \\ &= 37.2587629 + 62.7412371 \\ &= 100.00 \end{aligned}$$

Note that in the bond markets, no \$ sign would be associated with the above price. Bonds are normally issued in multiples of \$1,000 for corporate bonds and \$10,000 for government bonds in the US. To keep things consistent, bond prices are

frequently calculated as a % of face value. Thus, the *price* means that the bond sells at that percentage of its face value. To convert to the dollar price, simply multiply the FV of the bond by the *price* (% of its face value). Because of the various face values for bonds in the US, it is important to keep calculation to at least seven decimal places for accuracy.

From the example above, we see that a bond priced on its coupon date at the same yield as its own coupon, will be worth only its face value—no more and no less.

### Example 2.10

Suppose that the market yield is in fact 13% for the bond in Example 2.9. Assume the same 12% coupon and 4 years term left to maturity. Then the price for the bond will be:

$$\begin{aligned}
 P &= 6 \times PVIFA_{8,.065} + \frac{100}{(1+.065)^8} \\
 &= 6 \times \frac{1 - (1.065)^{-8}}{.065} + \frac{100}{(1 + 0.065)^8} \\
 &= 36.5325058 + 60.4231188 \\
 &= 96.9556245
 \end{aligned}$$

In this example, the price of the bond is given to seven decimal places. This is necessary as the majority of such bonds are bought and sold in amounts of ten million or more.

When market yield increases above the coupon rate, the price of the bond will fall below its face value—the bond is said to sell at a *discount*. On the other hand, if the yield were 11% instead of 12%, the price would be 103.167 2830. Since the bond in this case is selling above its face value, it is said to be selling at a *premium*. This will occur whenever the market yield falls below the coupon rate.

Formulas (9) and (10) assume that the bond is valued with one full discounting period remaining before the first coupon is received, such that  $n$  is a whole number. What happens when  $n$  is not a whole number?



**Sigma notation**

Recall Equation (9), which we reproduce below:

$$P = \frac{C}{(1+i)} + \frac{C}{(1+i)^2} + \dots + \frac{(C + FV)}{(1+i)^n}, \quad (9)$$


Formula (9) can also be written:

$$P = \sum_{k=1}^{k=n} \frac{C}{(1+i)^k} + \frac{FV}{(1+i)^n}. \quad (9')$$

The sigma sign refers to a sum ( $\Sigma$  is the Greek S). The representative index  $k$  in the sum is specified to range from 1 to  $n$ , reproducing the terms involving  $C$  in Equation (9).

**Bond pricing —less than one discounting period from the next coupon**

Equation (9) can be modified to deal with fractional time periods. To illustrate this, suppose that coupons are paid half yearly, the unit of time is half a year and the market yield  $i$  is also on a six-month basis (half the annual rate). Figure 2.2 is the relevant timeline. Inside a given half-year there are  $b$  days (so  $b$  may equal 180, 181, 182, 183, or 184 days according to the coupon month and the day-count convention). You want a price with ' $a$ ' days to the next coupon date.



How would you value the bond at the required point in time? Well, we know how to find the value  $PV_*$  at the next coupon date. In terms of sigma notation (see Equation (9') above) it would be:

$$PV_* = \left( \left( \sum_{k=0}^{k=n} \frac{C}{(1+i)^k} \right) + \frac{FV}{(1+i)^n} \right).$$

Using fractional discounting, the value at the required point must be:

$$P = \frac{1}{(1+i)^{\frac{a}{b}}} \times PV_*$$

Combining the two steps, this turns out to be same as writing:

$$P = \left( \left( \sum_{k=0}^{k=n} \frac{C}{(1+i)^{\frac{a}{b}+k}} \right) + \frac{FV}{(1+i)^{\frac{a}{b}+n}} \right), \quad (11)$$

which is the general formula for the bond price with fractional discounting.

Fractional discounting assumes that the periodic interest rate applies to the fractional time period. Be careful when using fractional discounting—it's only valid when the financial contract stipulates it in writing, or when market participants use it by convention. In the case of bond pricing, markets adopt this form of fractional discounting by convention.

Here's an additional wrinkle. There are two prices for bonds that sell in the United States (and other countries): the invoice price and the quoted price. The *invoice price* is the P from Equation (11). This is the total that you will pay for the bond. The *quoted price* will be lower by the amount of interest that has accrued since the last coupon payment. This interest amount is not a compounded amount, which is what we are doing in Equation (11), but is a simple interest amount. Mixing two interest payment systems—simple and compounding—is a curiosity of bond math which tends to be a challenge for the novice price maker.

The quoted price formula is,

$$P_Q = \left( \sum_{k=0}^{k=n} \frac{C}{(1+i)^{\frac{a}{b}+k}} + \frac{FV}{(1+i)^{\frac{a}{b}+n}} \right) - C \times \left( 1 - \frac{a}{b} \right) \quad (12)$$

Transforming Equation (11) into a closed form gives us:

$$P = \left[ (1+i)^{-a/b} \right] \times \left[ C \times (1 + PVIFA_{i,n}) + FV \times PVIF_{i,n} \right] . \quad (13)$$

Transforming Equation (12) into a closed form gives us,

$$P_Q = \left[ (1+i)^{-a/b} \right] \times \left[ C \times (1 + PVIFA_{i,n}) + FV \times PVIF_{i,n} \right] - C \times \left( 1 - \frac{a}{b} \right). \quad (14)$$

To make matters worse,  $a$  and  $b$  are not the same for all bonds. For US federal agency and corporate bonds,  $a$  is calculated using a thirty day month and  $b$  is based on a 360 day year (180 day half year). For US Treasuries,  $a$  is the actual number of days from settlement date to the next interest date, and  $b$  is the actual number of days between coupons.

### Example 2.11

Find the invoice price for the March 31, 2003 US Treasury issue with a coupon rate of 4.25% p.a. The market yield is currently 4.283% and settlement date is April 19, 2001. The semi-annual coupons are \$2.125 per \$100 face value ( $C = 4.25$ ). The number of full coupon periods left to run is three ( $n = 3$ , from September 30, 2001 to March 31, 2003). The number of days from settlement date to the next interest date is 164 days ( $a = 164$ ). The number of days from the last interest date, March 31, 2001, to the next interest date, September 30, 2001, is 183 days ( $b = 183$ ). The current market yield is 2.1415% per half year ( $i = 0.021415$ ).

We will use Equation (13) for pricing the bond. Per \$100 face value, the bond will price at:

$$\begin{aligned} P &= \left[ (1 + .021415)^{\frac{-164}{183}} \right] \times [2.125 \times (1 + 2.875953) + 100 \times .9384115] \\ &= 0.9811902 \times [102.0775468] \\ &= 100.1574861 \end{aligned}$$

This determines the total amount to be paid for the bond – the invoice amount.

The quoted price per \$100 face value is the invoice price less the accrued interest. The term we have to subtract is:

$$C \times \left( 1 - \frac{a}{b} \right) = 2.125 \times \left( 1 - \frac{164}{183} \right) = 0.2206284$$

Hence, the quoted price per \$100 face value is:

$$100.1574861 - 0.2206284 = 99.9368577 \cong 99.94$$

### Zero-Coupon Bonds

Zero-coupon bonds consist of just a single cash flow, paid at the bond's maturity date. There are no intervening coupons. In the US, taxes are computed on capital gain as well as ordinary income. This places zero-coupon bonds at a psychological disadvantage to coupon bonds. People do not like paying taxes on something from which they aren't earning income. This is why zeros are popular in IRAs and other tax advantaged accounts. In some countries where tax is assessed on a cash flow basis, these zeros enjoy a considerable tax advantage. In addition, zero coupon bonds have considerable theoretical importance in finance.

In general, zero coupon bonds:

- have an initial term of greater than 1 year
- have no intermediate cash flows
- are always purchased below face value
- usually have six-month compounding.

The zero pricing equation is:

$$P = \frac{FV}{\left(1 + \frac{r}{2}\right)^n} \quad (15)$$

where  $r$  = the nominal annual zero-coupon rate expressed as a decimal and  $n$  = the number of half-years to maturity.

#### Example 2.12

Price a ten year zero-coupon bond with face value \$1m at a market yield of 7% p.a. Assume the usual semi-annual compounding.

$$P = \frac{\$1,000,000}{\left(1 + \frac{07}{2}\right)^{20}} = \$502,565.88$$

If the market yield increases to 8%, the bond re-prices to:

$$P = \frac{\$1,000,000}{\left(1 + \frac{.08}{2}\right)^{20}} = \$456,386.95 .$$

The bondholder will incur a capital loss of \$46,178.93, a very significant “hit” for an initial investment of \$502,565.88. We can express this as a percentage loss:

$$\frac{\$46,178.93}{\$502,565.88} \times 100\% = 9.18863\% .$$

Note that this capital loss will not be realized until the holder sells the bond. It does, however, represent the *opportunity cost* of holding the instrument.

The problem of potential capital loss is an extremely important one for those who hold investments or issue liabilities, so it is worthwhile to compare the sensitivity of such an instrument to that of other financial instruments. Such information will be highly useful to portfolio managers making decisions about purchases and sales (or issues and buy backs) of assets (liabilities) given various expected interest rate changes.

To illustrate this, consider a much shorter-term security, for example, a discount instrument such as a CD, and evaluate the change in its value given a change in interest rates.

### Example 2.13

Price a 90 day CD with a face value \$1m when the market yield is 7%. Recall that we can price this security directly using Formula (4) from Chapter 1, because CDs are priced using the bond equivalent approach:

$$Price = \frac{FV}{1 + \left(i \times \frac{n}{365}\right)} .$$

where:

FV = face or nominal value of the security

$n$  = number of days from settlement date to maturity

$i$  = (nominal annual) yield for securities of  $n$  days to maturity.

Given the numbers quoted above, we have:

$$Price = \frac{\$1,000,000}{1 + \left(0.07 \times \frac{90}{365}\right)} = \$983,032.59 .$$

Now, suppose that the instant after you bought it, the market re-prices the 90-day CD yield at 8% p.a. The new price is:

$$Price = \frac{\$1,000,000}{1 + \left(0.08 \times \frac{90}{365}\right)} = \$980,655.56 .$$

If you decide to sell with 90 days still to run, you will record a capital loss of \$2,377.03. On the other hand, if the CD is held to maturity you will earn 7% instead of 8%, which could have been earned had the purchase been better timed. You have incurred an opportunity loss of 1%.

Expressing this capital loss as a percentage:

$$\frac{\$2,377.03}{\$983,032.59} \times 100\% = 0.241806\% .$$

In summary, when interest rates move from 7% to 8%, there is a capital loss of 0.24%.

Clearly this instrument is far less price sensitive to changes in the interest rate than a much longer term instrument such as the ten-year zero coupon bond in Example 2.10, where the capital loss was 9.2%. The issue of value sensitivity is going to be important for fixed interest managers, and we take this up in the next section.

## 2.7 Duration and value sensitivity

While the concept of duration is commonly applied to bonds and other traded fixed interest instruments, it can be applied to any set of cash flows. Duration is a measure of the *average weighted times to receipt* for a given set of cash flows.

For example, suppose we are to receive \$1000 in one month plus \$1000 in three month's time. We could say that on average we are to receive \$2000 in two months. The simple average term of this cash flow would be approximately two months.

Duration is measured however, not as the simple average of the timing of payments (as for the cash flows above), but as a *weighted average of the timing of payments* where the weights are the present values of the cash flows. This makes allowance for the time value of money.

Suppose we have a series of cash flows, accruing at times  $t_i$  into the future, where  $i$  ranges from  $i=1$  to  $i=n$ . In other words, cash flow #1 occurs at time  $t_1$ , cash flow #2 at time  $t_2$ , and so on. Let cash flow # $i$  have present value  $PV_i$ , and let PV.

$\sum_{i=1}^n t_i PV_i$ , be the total present value of the set of cash flows. Duration is defined by:

$$\begin{aligned}
 D &= \frac{\sum_{i=1}^n t_i PV_i}{\sum_{i=1}^n PV_i} \\
 &= \sum_{i=1}^n t_i \left( \frac{PV_i}{PV} \right)
 \end{aligned} \tag{16}$$

Version (16) makes it clear that each time to receipt  $t_i$  is weighted by its share of the total PV generated.

In the case of a bond,  $PV = P$ , the bond's market value, and the  $PV_i$  refer to the cash flows generated by the coupons and the final return of principal.

If the PV's are calculated on a yield to maturity basis (as we have so far assumed in this chapter), the resulting  $D$  is called *Macaulay duration*. Another version, called *Fisher Weil duration*, utilizes zero coupon rates to calculate the  $PV_i$  (See Chapter 3).

**Example 2.14**

We will calculate Macaulay duration for a 10% coupon bond, which has two years left to run and a yield to maturity of 8.0%. Assume a face value of \$1m. We can use a spreadsheet to generate a table such as the one that follows.

**Table 2.2 Calculation of Macaulay Duration for two-year bond**

Time ( $t_i$ )	Cash flow	Present Value $PV_i$	$t_i * PV_i$
0.5	50,000	$\frac{50,000}{(1.04)} = 48,076.92$	24,038.46
1.0	50,000	$\frac{50,000}{(1.04)^2} = 46,227.81$	46,227.81
1.5	50,000	$\frac{50,000}{(1.04)^3} = 44,449.82$	66,674.73
2.0	1,050,000	$\frac{1,050,000}{(1.04)^4} = 897,544.40$	1,795,088.80
Totals		PV = 1,036,298.95	1,932,029.80

To simplify the spreadsheet calculation, we have designated  $t$  in years so that  $D$  is also expressed in years. (If we had calculated this using six-month periods, the result would have been twice as large, or 3.7287113 six-month periods). We have:

$$D = \frac{1,932,029.80}{1,036,298.95}$$

$$= 1.864 \text{ years}$$

Note that the duration of the bond is less than its term to maturity, which is two years.

In the case of a bond, Equation (16) has a closed form version.

$$D = \frac{1+i}{i} - \frac{(1+i) + n(c-i)}{c[(1+i)^n - 1] + i} \quad (17)$$



Where:  $i$  = the effective interest rate for the period  
 $c$  = the coupon rate for the period expressed as a decimal  
 $n$  = the number of compounding periods

Let's continue with Example 2.14 by calculating the duration, using Formula (17).

$$\begin{aligned} D &= \frac{1+.04}{.04} - \frac{(1+.04) + 4 \times (.05 - .04)}{.05 * [(1.04)^4 - 1] + .04} \\ &= 26.0 - 22.2712887 \\ &= 3.7287 \text{ six month periods or 1.864 years} \end{aligned}$$

For a bond of maturity  $n$  periods, notice that if the coupon rate is equal to zero, the formula collapses to simply  $n$ . Therefore, the duration of a zero coupon bond is simply equal to its time to maturity.

Not quite as obvious is the closed form solution for the duration of an annuity. With a little more algebraic manipulation we can get the duration for an annuity of  $n$  periods as equal to:

$$D = \frac{1+i}{i} - \frac{n}{[(1+i)^n - 1]} \quad (18)$$

Duration is employed in the measurement and management of interest rate risk.

Duration is used as a *measure of the price sensitivity* of an instrument (or that of a portfolio of instruments) to changes in interest rates. This use results from the relationship between the instrument's duration and the percentage change in the instrument's price caused by a change in interest rates.

We can describe the sensitivity relationship numerically as

$$\frac{\Delta P}{P} \approx -D \times \left( \frac{\Delta i}{1+i} \right) \quad (19)$$

where:  $\Delta$  represents a small change  
 $P$  is the price of the instrument  
 $D$  is the duration in interest periods  
 $i$  is the (effective) period interest rate

The left-hand side of Equation (19) is the percentage change in the price of the instrument. The right hand side of the equation indicates that  $D$  is approximately the coefficient of changes in the interest rate as they affect price. This equation frequently appears in a slightly different form as:

$$\frac{\Delta P}{P} \approx -D' \times \Delta i \quad (19')$$

where:  $D' = D / (1 + i)$  and is known as *modified duration*.

For Example 2.14 above we can calculate modified duration as:

$$D' = \frac{3.7287113}{1 + \frac{0.08}{2}} = 3.585 \text{ periods} = 1.79 \text{ years}$$

The interpretation is that a 100 basis point increase in the interest rate (for example, from  $i=10\%$  to  $i=11\%$ ), will produce a capital loss of 1.79%.

Another use of duration is in the area of *portfolio immunization*. This approach allows the investor to eliminate exposure to interest rate risk by matching the duration of an investment portfolio to the required investment horizon.

To see how this works, suppose you had a commitment to pay \$1 million in five years time. You could buy a five-year zero coupon bond with a face value of \$1 million. Now suppose five years have passed and interest rates have changed. Would this affect your ability to meet the contract? Clearly not – your bond will pay exactly \$1 million now and you are OK.

But would this be true if instead of a five-year zero coupon, you had bought—say—a ten-year coupon bond, intending to sell it after the five years? Suppose that near the end of the fifth year, interest rates had risen. Would you then be able to meet your \$1 million contract? This is doubtful, as the price of the bond would be lower.

Instead, you can ensure that whatever bond or portfolio you use, its initial duration is set equal to the investment horizon (in this case, five years). This effectively turns your portfolio or coupon bond (in terms of interest rate dynamics) into an

equivalent zero coupon bond. You can make this exact by a small amount of rebalancing as time passes, to keep the duration equal to the remaining investment horizon. For instance, you would occasionally need to swap the current investment for a coupon bond of slightly lower duration. But for most horizons, the rebalancing required is small.

## 2.8 Interest rate futures

Forward contracts are agreements calling for future delivery of an asset at a price agreed upon at the inception of the contract. A futures contract is a more tradable form of a forward contract that is marked to market each day, with the change in value being credited or debited to your account with a broker. To ensure tradability, futures contracts are written on underlying assets (physicals) of standardized units, quality and quantity. Futures contracts are traded on exchanges, while forward contracts are not—they are OTC (“over the counter”) instruments. In this chapter, we limit our discussion to interest rate futures contracts.

At least 30 interest rate futures contracts are traded on the Chicago Board of Trade (CBOT) and the Chicago Mercantile Exchange (CME). Summary listings of these contracts are provided in Tables 2.4 – 2.6 in the appendix to this chapter.

Futures are used for two main purposes—hedging and taking open positions (speculating). Some people need to reduce or eliminate price movements in the underlying commodity. The mitigation of such natural exposures is known as hedging. Others try to profit from price movements in the underlying commodity. Because you need only put down between 5% and 15% margin, depending on the contract, futures are heavily leveraged instruments favored amongst speculators.

The use of derivatives in risk management and speculation is beyond the scope of this book. However, we will look briefly at three examples of future contracts—the T-bill, the Eurodollar and T-bond contracts—which are used for the purpose of hedging interest rate risk.

### U.S. T-Bill futures

The 13-week (three month) Treasury bill contract is based on a simple discount security, the T-bill that was discussed in Section 1.3 of Chapter 1.

**Example 2.15****Table 2.3**

Open	High	Low	Settle	Change	Discount Yield	Discount Change	Open Interest
93.69	93.69	93.64	93.64	-0.09	6.36	0.09	221

Table 2.3 captures the “feel” of how 13-week T-bill futures would be presented in the *Wall Street Journal*. The only difference is that there would be nothing indicating that the yield is a discount yield – the *WSJ* assumes that the reader knows this. The open, high, low and settlement price are not “true” prices, but are another way of specifying the annual discount interest rate. A settle price of 93.64 means an annualized discount from par of 6.36% ( $100\% - 93.64\% = 6.36\%$ ). This 6.36% is an annual rate based on 90 days and a 360-day year. Converting it to a 90-day discount yield is done by multiplying  $90/360$  by  $6.36\% = 1.59\%$ .

If you buy this contract, the invoice price you are promising to pay is:

$$P = \$1,000,000 \times \left( 1 - .0636 \times \frac{90}{360} \right) = \$984,100.00$$

Notice that P is not subscripted with a 0. The reason is that technically this transaction will take place in the future at the contract expiration date, though most are closed out prior to this date at the market price quoted at the time. The latter will be listed next to the T-bill contract in the *WSJ*. For a complete listing of contract expiration dates and other information for contracts traded at the Chicago Mercantile Exchange, go to <http://www.cme.com/clearing/spex/XMLReports/intrRateGroup.htm>.

What happens if the interest rate (discount) drops by one basis point to 6.35%? The price for a new contract for will now be:

$$P = \$1,000,000 \times \left( 1 - .0635 \times \frac{90}{360} \right) = \$984,125.00$$

Notice the difference between the original contract and the new one is \$25.00, which is exactly two (half basis point) ticks for this contract (see Table 2.4 in the appendix to this chapter).

If the interest rate rises by one basis point the price of a new contract will be:

$$P = \$1,000,000 \times \left( 1 - .0637 \times \frac{90}{360} \right) = \$984,075.00$$

This time the difference in value between the two contracts is -\$25.00—once again, exactly two (half basis point) ticks. In each case, the gains and losses will be posted to your broker account. Where current losses eat too much into the margin you have lodged with the broker, you will be required to front up more cash to restore the margin.

To evaluate gains or losses when trading this contract, it isn't actually necessary to go through all the calculations that we performed above. We can measure such gains and losses for the T-bill contract simply by

### The Eurodollar time deposit contract

Another three-month contract trading on the CME—one that enjoys greater popularity and liquidity than the T-bill contract—is the Eurodollar Time Deposit contract. It has a principal value of \$1,000,000 and trades in one-tick (= one basis point) movements worth \$25.00 per basis point for all but the current month contract, for which the contract trades in half-tick movements (= half basis point) worth \$12.50 per half-tick. Like the T-bill contract, the Eurodollar contract is quoted using a price index, which is derived by subtracting the current interest rate for the contract from 100.00. For example, an interest rate of 4.50% translates to a futures index price of 95.50 (100.00 - 4.50 = 95.50).

To measure gains or losses on Eurodollar futures positions, the calculation is:

- Number of ticks moved  $\times$  \$25.00  $\times$  number of contracts held.

For the current month contract this calculation becomes:

- Number of half-ticks moved  $\times$  \$12.50  $\times$  number of contracts held.

Eurodollar positions that are not closed out prior to settlement date are cash settled using a final settlement price based upon the British Bankers Association (BBA) rate fixing. The rate that applies is the BBA Interest Settlement Rate for three-month U.S. dollar deposits determined at 11:00 a.m. London time on the contract's last trading day. For Eurodollar contracts held until contract closeout, the final settlement price is rounded to four decimal places.

### Example 2.16

In late June, a corporate treasurer reviewing his company's cash flow projections, estimates borrowing requirements of \$10,000,000 for mid-September. The floating-rate bank loan the treasurer has negotiated will be priced at one percent over the three-month Eurodollar LIBOR rate at the time of funding (with interest paid quarterly). LIBOR is currently quoted at 5.35%, while September Eurodollar futures, which can be used to lock in a September borrowing rate, are trading at 94.45, giving an implied forward rate of 5.55% (calculated as  $100.00 - 94.45$ ).

The corporate treasurer, concerned that rates will rise over the time period prior to drawing down the loan, decides to "fix" the borrowing rate by selling ten Eurodollar contracts. Since the September Eurodollar contract will trade very close to the three-month Eurodollar LIBOR rate by mid-September, the borrowing rate achieved by the treasurer will be very close to 6.55% (the Eurodollar futures rate of 5.55% plus the bank's 1% spread over LIBOR).

The hedge works in the following way:

On June 29, the corporate treasurer sells ten Eurodollar futures contracts at 94.45 (i.e. a rate of 5.55%). Suppose that by September 15<sup>th</sup>, the Eurodollar LIBOR physical rate has risen to 5.75%. The loan rate, including the bank's margin, will be at 6.75%. However, if the corporate treasurer is able to buy ten Eurodollar futures contracts at the same LIBOR interest rate (and it will be very close if not exactly the same), the futures position will be closed out at 94.25 for a gain of 20 basis points (94.45-94.25). The 20 basis points earned on the futures contracts will effectively reduce the cost of the loan from 6.75% to 6.55%.

### The 30-year T-bond futures

Let us briefly turn our attention to another popular interest rate futures contract—the 30-year Treasury bond. Because of the complexity of this contract, we will describe only its important features. If you are interested in more detail, visit the Chicago Board of Trade (CBOT) website at: <http://www.cbot.com>.

The underlying instrument is a hypothetical 20-year 6% semi-annual coupon Treasury bond. The quote in the *Wall Street Journal* will be three digits followed by a hyphen and two more digits. For example, if you see 101-26 it means that the contract is trading at 101 and 26/32's of the \$100,000 par value. This contract is not a cash-settled contract. If you hold the instrument until maturity you must deliver (if you shorted the contract) or buy (if you are long the contract) one T-bond for every contract.

If you are the one who is delivering the bond, what happens if you can't find a 20-year T-bond with a 6% semi-annual coupon? Fortunately, you can deliver a T-bond that has a minimum of 15 years before the Treasury can call it. Suppose you have found a 5% 20-year bond, which the CBOT has deemed acceptable for satisfying the contract. But the value of this bond is less than the value of the underlying bond of the futures contract. Receiving this would be unfair to the other party. On the other hand, if all you can find is a 7% 20-year bond that is more valuable now, this would be unfair to you. If these issues went unresolved there would be little demand for these contracts.

To resolve these issues, the CBOT uses conversion factors. The conversion factor (CF) is calculated:

$$CF = \sum_{t=1}^n \frac{cashflow_t}{1.03^t} / 100,000$$

Conversion factors are updated and published by the board of trade in Excel and downloadable by going to <http://www.cbot.com/cbot/docs/22879.xls>. These factors are constant throughout the trading period of the contract. The invoice price paid by the buyer of the Treasury bonds delivered to the seller is calculated by the following equation:

Invoice price = contract size × futures contract × settlement price conversion factor

There is much more to be said about the 30-year T-Bond futures contract and futures contracts in general. The Chicago Board of Trade and the Chicago Mercantile have excellent websites with suggested reading for those who wish to explore this area in more detail.

**Example 2.17**

Suppose that the contract settles at 97-00, which is 97% of par. The conversion factor for the bond being delivered is 1.05 (just downloaded from the CBOT website). Thus, the invoice price of the bond to be delivered is:

$$\$100,000 \times .97 \times 1.05 = \$101,850$$



## Interest Rate Futures

**Table 2.4 Interest Rate Futures**  
Exchange: Mercantile

Underlying Instrument	Ticker Symbol	Contract Size	Tick Size	Delivery
Eurodollar	ED	Eurodollar Time Deposit having a principal value of \$1,000,000 with a three-month maturity	1 point = .01 = \$25.00 per contract	Cash
13 Week T-Bill	TB	3-month (13-week) U.S. Treasury Bills having a face value at maturity of \$1,000,000	½ point = .005 = \$12.50	Cash
LIBOR FUTURES	EM	Eurodollar Time Deposit having a principal value of \$3,000,000 with a one-month maturity	½ point = .005 = \$12.50	Cash
FED FUNDS TURN FUTURES	TZ	\$45,000,000 overnight Fed Funds deposit	1 point = .01 = \$12.50	Cash
10 YR AGENCY FUTURES	F0	\$100,000 face value of the 10 year Fannie Mae Benchmark Notes or Freddie Mac Reference Notes	1 point = .01 = \$15.625	Cash
5 YR AGENCY FUTURES	F5	\$100,000 face value of the 5 year Fannie Mae Benchmark Notes or Freddie Mac Reference Notes	1 point = .01 = \$15.625	Cash
ARGENTINE 2X FRB BRADY BOND FUTURES	AT	\$100,000 Current original face value.	1 point = 0.01 = \$10.00	Cash
ARGENTINE PAR BOND FUTURES	AX	\$100,000 Current original face value.	1 point = 0.01 = \$10.00	Cash
BRAZILIAN 2X C BRADY BOND FUTURES	BF	\$100,000 Current original face value.	1 point = 0.01 = \$10.00	Cash
BRAZILIAN 2X EI BRADY BOND FUTURES	BE	\$100,000 Current original face value.	1 point = 0.01 = \$10.00	Cash
MEXICAN 2X BRADY BOND FUTURES	MN	\$250,000 of the original face value of the Mexican Par Brady Bond with warrants	1 point = .01 = \$25.00	Cash
EURO YEN FUTURES	EY	Euro yen Time deposit having a principal value of 100,000,000 Japanese yen with a three-month maturity	½ point = .005 = 1,250 Yen	Cash
JGB FUTURES	JB	10,000,000 Yen	1 point = .01 = 1,000 Yen	Cash
EURO YEN LIBOR FUTURES	EL	Euro yen Time deposit having a principal value of 100,000,000 Japanese yen with a three-month maturity	½ Point = .005 = 1,250 Yen	Cash
MEXICAN TIE FUTURES	TE	Mexican interbank loans of 28-day maturity having a principal value of 6,000,000 Mexican pesos.	1 point = .01 = 50.00 Mexican pesos	Cash
MEXICAN CETES FUTURES	TS	Mexican Treasury Bills of 91-day maturity having a principal value of 2,000,000 Mexican pesos.	1 point = .01 = 50.00 Mexican Pesos	Cash

Contracts on the Mercantile Exchange are settled in cash, unlike larger contracts on the Chicago Board of Trade.

**Table 2.5 Interest Rate Futures**  
Exchange: Chicago Board of Trade

Underlying Instrument	Ticker Symbol	Contract Size	Tick Size	Delivery
Current Mortgage Price Index	MF	\$1,000 times Index for a par of \$100,000	1/32 of one pt. Or \$7.8125 rounded to nearest cent	Class A Mortgage backed securities
Long Term Municipal Bond Index Futures	MB	\$1,000 times The Bond Buyer™ 40 Index	1/32 of one pt. or \$31.25/contract	Cash on last trading day of the month
5 Yr US Treasury Notes	FV	One US Treasury Note with a par of \$100,000	1/32 of one pt. or \$31.25/contract	5 Yr US Treasury Note
5 Yr Agency Notes	DF	One Fannie Mae Benchmark Note or Reference Note with a par of \$100,000	½ of 1/32 of one pt. or \$15.625/ Contract	One Fannie Mae Benchmark Note or Reference Note
30 Yr US Treasury Bonds	US	US Treasury Bond with a par of \$100,000 or multiple	1/32 of one pt. or \$31.25/contract or multiple	US Treasury Bond with a par of \$100,000 or multiple
30 Day Federal Funds	FF	\$5m	\$20.835 per ½ of one basis point rounded to nearest cent	Cash settled against the average daily fed funds rate for the delivery month
2 Yr Treasury Notes	TU	US Treasury Notes with a par value of \$200,000 or multiple	1/4 of 1/32 of one pt. or \$15.625/ Contract	US Treasury Notes with a par value of \$200,000 or multiple
10 Yr Treasury Notes	TY	US Treasury Note with a par of \$100,000 or multiple	½ of 1/32 of one pt. or \$15.625/ Contract	US Treasury Note with a par of \$100,000 or multiple
10 Yr Agency Notes	DN	One Fannie Mae Benchmark Note or Reference Note with a par of \$100,000	½ of 1/32 of one pt. or \$15.625/ Contract	One Fannie Mae Benchmark Note or Reference Note with a par of \$100,000

The Chicago Board Of Trade offers smaller contracts that are settled in cash. The Mid American Exchange formerly offered these contracts, but this exchange is now part of CBOT.

**Table 2.6 Interest Rate Futures**  
Exchange: Chicago Board of Trade - Mid Am

Underlying Instrument	Ticker Symbol	Contract Size	Tick Size	Delivery
US Treasury Bonds	XB	\$50,000 par value	1/32 of one pt. Or \$15.62 per contract	Cash
US Treasury Bills	XT	\$500,000 par value	1/2 basis point or \$6.25/contract	Cash
Eurodollars	UD	\$500,000	1 basis point or \$12.50/contract	Cash
5 Yr Treasury Notes	XV	\$50,000 par value	1/32 of one pt. or \$7.81/contract	Cash
10 Yr Treasury Notes	XN	\$50,000 par value	½ of 1/32 of one pt. or \$15.625/Contract	Cash

1. An estate pays an annuity certain (cash payable whether or not the annuitant survives) of \$50,000 per annum for five years. Assume the payments are made at the end of the period. From the sixth year until the tenth year, this annual payment will double. Determine the present value of the estate assuming 8% per annum compound.
2. Find the present value of \$500 payable monthly for six months. The first of the six payments occurs five months from now. The compound interest rate is 1.1% per month.
3. Find the effective rate p.a. if the nominal rate per annum compounded quarterly is 12%.
4. 12% p.a. nominal compounded monthly is equivalent to what nominal rate p.a. compounded quarterly?
5. 8% p.a. nominal compounded half-yearly is equivalent to what nominal rate p.a. compounded monthly?
6. Consider a bond with maturity date March 1, 2009, coupon 7% p.a. (paid semi-annually). The market yield increases from 7% to 8% on March 1, 2001. On a face value of \$1m., what is the percentage change in price? Is this instrument as sensitive as a zero coupon bond of identical final maturity priced on the same day as the rate moves from 7% to 8%?
7. Price the coupon bond of Question 6 on the March 2, 2001 at 8% market yield. Why is there a price difference? (This problem is more difficult as there is a partial coupon period).
8. Price a seven-year zero coupon bond with face value \$1m at a market yield of 8% p.a. Assume semi-annual compounding. Price the same bond at a market yield of 7% p.a. and compare its sensitivity with that of the bonds of Question 6.
9. Calculate Macaulay duration for the bonds in Question 8, first at a yield of 7%, then at 8%.
10. Price a 180 day T-bill with face value \$1m when the bank discount rate is 7%.

11. Price a 180-day CD at a market yield of 7%. Why is the price different from that of the T-Bill in problem 10?
12. Suppose that the table below lists prices of two Eurodollar time deposit contracts quoted on the Chicago Mercantile Exchange on June 12, 2001.

Futures Contract	BID	OFFER
SEP 01 ED	94.25	94.26
DEC 01 ED	94.16	94.17

The date is June 12, 2001. A corporate treasurer wishes to increase working capital \$40 million by drawing down a bank loan based on six-month LIBOR on September 14, 2001. What financial risk does the corporate treasury face now, and how can this risk be hedged using Eurodollar futures? Describe the hedge stating how many of which contract(s), the transaction involved (buy or sell) and the likely dealing price.

13. On June 12, 2001 a corporate treasurer managing a liability portfolio consisting of 22 million of ten-year fixed rate debt is considering interest rate forecasts in the ten-year area of the yield curve.

To what risk is the portfolio exposed, and why? How could the treasurer cover against this exposure using future contracts? Select the most appropriate contract from Table 2.3. Then describe the hedge stating which contract, approximately how many to use, and the transaction involved (buy or sell). Would such a hedge remove all interest rate risk for the debt portfolio?



---

# Three

## Zeros, Forwards, and The Term Structure

Until now, the notion of time and the way it is embedded into interest rates have been treated in a simplistic fashion. You may have noticed that although we were dealing with the valuation of cash flows of perhaps quite long maturity structures, the near-term parts of those cash flows were valued in the same way as the long-term parts, using the same interest rate as a discount factor. You might have wondered whether this is a good assumption to make, and you'd be right to wonder. Although it will suffice for an initial approach – and a good deal of practical valuation is done this way– it is not theoretically the best way to proceed. So we must take a closer look at this issue. Even if you never end up using the tools of this chapter, reading through it will reinforce your understanding of the concept of interest rates.

This serves to introduce the topics of this chapter: zero coupon rates, implied forward rates, and the way that they re-assemble into the term structure of interest rates. The term structure is all about the dependence of the interest rate structure on forward, or maturity, time. The notion of time itself gets a bit tricky, so we might do well to start with some clarifying remarks on what we mean by time in this context.

### 3.1 Notation and definitions

Arbitrage	A zero-risk, costless investment strategy that generates a profit. A no-arbitrage pricing or equilibrium is where no such opportunities exist.
Long	A position that benefits from a price rise and loses from a price fall. In portfolio terms, a positive holding.
Short	A position that benefits from a price fall and loses from a price rise. In portfolio terms, a negative holding.
Hedge	An instrument or position taken to reduce the exposure of the portfolio to adverse changes in value. A perfect hedge reduces the exposure to zero, but typically also results in zero upside potential.
Unit-discount bond	A hypothetical zero coupon bond with a notional \$1 face value.
Zero coupon bond	Debt instrument that has no coupons, but merely a promise to pay the face value at maturity. The reward element derives from the discount at which it sells, relative to its face value.
Zero coupon rate	The yield to maturity of a zero coupon bond; varies with maturity time.
$R_T$	The yield to maturity or IRR on a zero coupon bond of maturity $T$ .
$f_t$	The implied forward rate is the implied one period interest rate for future period $t$ priced into an instrument such as a bond or swap by the market at current time.
Term structure of interest rates	A graph of zero coupon rates, yields, or implied forward rates against maturity time.
Yield curve	The term structure curve applied to yields.
Spot rate	The interest rate for the one period immediately ahead, known as of the current date.



**Bootstrapping**      A sequential methodology using actual bond prices and coupons to derive the zero coupon term structure.

### 3.2 A brief mystery of time

Suppose you are at time 0 ( $t = 0$ ), the present time, and are looking at T-bill rates. You are concerned about them, because you know that your very expensive adjustable-rate mortgage is partially tied to them.<sup>1</sup> You might form an impression of what the T-bill rate will be for every six-month period throughout the next ten years, as though you had a mental image of the rate at six-month intervals from now ( $t = 0$ ) until  $t = 20$ . For convenience, we have adopted the basic time interval of six months.

The notion of time for this mortgage planning scenario is both a current as well as a forward concept. The expression *forward time* (and *forward rate* in the next paragraph) is adopted to distinguish it from the current time (and the current rate, which will be known as the *spot rate*). Thus, forward time (i.e.  $t > 0$ ) must always have the current time as a base. Here, forward time is the 20 six-month periods following current time,  $t = 0$ . Current time is in *real time*, while *forward time* refers to the future. Real time refers to the actual passage of time.

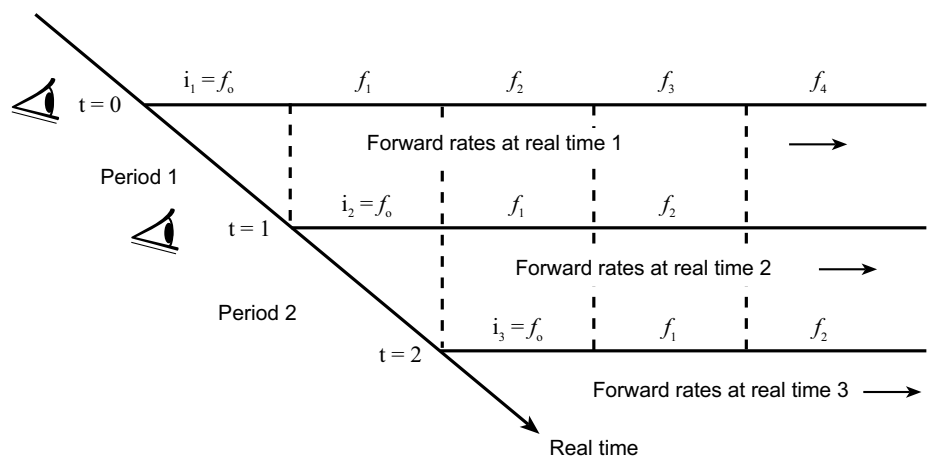
Suppose that at real time  $t = 0$ , the T-bill rate over the coming six-month period is  $i_1$  (the periodic rate used in Chapters 1 and 2). You could also form an estimate of the next six-month T-bill rate, which is known as the forward interest rate  $f_1$ . It has to be a forward rate, because you won't know the true six-month T-bill rate at time 1 until six months has passed.

Figure 3.1 illustrates this concept. The eye represents you looking forward into the future. At any point in actual or real time (denoted  $t$ ), you will know what the current one period rate is, because this is set by the market now to apply over the coming unit period (here we use six months). At real time  $t = 0$  you know  $i_1$ , at real time  $t = 1$  you know  $i_2$ , and so forth.

But at time  $t = 0$ , you don't know what  $i_2$  will be. Only when real time has moved forward to  $t = 1$  will you know, so in the meantime, you have to form an estimate of  $i_2$ . This is denoted as  $f_1$ . Now imagine you are six months older and can see what the actual six-months rate is in real time  $t = 1$ . The new spot rate ( $i_2$ ) will almost certainly differ from your guess at  $t = 0$  about the forward rate,  $f_1$ . In other words, generally speaking,  $i_2 \neq f_1$ .

In the following discussion, we shall call the  $i$  rates the *spot rates*. They are the actual one-period rates observed in real time. Hence,  $i_t$  means the actual one-period rate that applies for period  $t$ , measured at the instant in real time  $t - 1$ . By period  $t$ , we mean the time period which begins at the instant in real time,  $t - 1$ , and runs to real time,  $t$ . Period 1, therefore, runs from the instant real time  $t = 0$  to the instant real time  $t = 1$ .

We shall also assume general units of time, which may correspond to a year, a six-month period, or whatever time unit applies. By using  $t$  as the time unit, we assume that this is the basic discounting or rest period, so that interest rates are measured using the same unit of time. For example, if the annual rate is 4 % for six-month T-bills, and if  $t$  refers to six-month intervals, then the spot interest rate is 2%.



**Figure 3.1 Real versus forward time**

### 3.3 Zero coupon rates

A current one-period unit discount bond would be an instrument that pays back \$1 at future time 1 ( $t = 1$ ). It is the simplest example of a zero coupon bond, one for which the face value is just one dollar. Suppose that the market prices this bond at  $P_1$ . Its yield,  $R_1$ , is then given by:

<sup>1</sup>Lenders base ARM rates on a variety of indexes. Among the most common are the rates on one-, three-, or five-year Treasury securities. Another common index is the national or regional average cost of funds to savings and loan associations. A few lenders use their own cost of funds, over which – unlike other indexes – they have some control. As a borrower, you should find out which index will be used, how often it changes, how it has behaved in the past, and where it is published.

$$P_1 = \frac{1}{1 + R_1} \quad , \quad \text{thus } R_1 = \frac{1}{P_1} - 1.$$

Note that the subscript 1 indicates that this is a one-period instrument, as seen from real time  $t = 0$ . The one-period zero coupon rate is also the IRR – the interest rate that equates the price to present value of the \$1 terminal cash flow. Since this is a one-period instrument, the IRR is also a one-period rate.

Suppose that at  $t = 0$ , we have available a two-period pure discount bond. Technically, this is now a zero coupon instrument because there is no coupon paid at time 1. Its price and yield have an inverse relationship, following Equation (8) from Chapter 1, so that:

$$P_2 = \frac{1}{(1 + R_2)^2} \quad ; \quad \text{thus } R_2 = \left( \frac{1}{P_2} \right)^{0.5} - 1.$$


Will  $R_2$  be the same as the one-period rate  $R_1$ ? Not necessarily. After all, with the two-period instrument we are tying our money up for two periods—we might expect some extra compensation since this reduces our liquidity.

We can continue this example for three time periods, four time periods, and so on, solving for the various yields. In general:

$$P_t = \frac{1}{(1 + R_t)^t} \quad ; \quad \text{thus } R_t = \left( \frac{1}{P_t} \right)^{\frac{1}{t}} - 1 \quad (1)$$

All the interest rates we have derived – the zero coupon rates – are computed at  $t = 0$ . Consequently, everything is seen from the same point in real time – the present. The plot between  $R_t$  (zero coupon rates) and  $t$ , with  $t$  on the horizontal axis (maturity), is called the *zero coupon rate term structure of interest rates*. Note that by convention, all rates are annualized in the graph.

Figure 3.2 illustrates a zero coupon rate term structure. This particular version has the rate increasing with maturity; it is upwardly sloping. An upward sloping term structure is known as a *normal* yield curve. Most theories on term structure suggest an upward slope, and this is often what we see in the “real” world. The opposite—a downward sloping curve—where the shorter rates are higher than the longer rates, is known as an *inverted* yield curve. This would not be considered a normal state of affairs; it is characteristic of periods of tight monetary control. Humped curves also surface from time to time, and some evidence suggests they presage a recession.



**Figure 3.2: Zero coupon term structure**

If we knew the zero coupon rate structure—in other words, the entire sequence  $(R_1, R_2, R_3, \dots, R_T)$ —we could use this data to value any set of cash flows into the future. An obvious example is a coupon bond of maturity  $T$  periods and face value  $FV$ . Its zero coupon based valuation would be:

$$P_T = \frac{C}{1+R_1} + \frac{C}{(1+R_2)^2} + \dots + \frac{(C+FV)}{(1+R_T)^T} \quad (2)$$

If market participants evaluate coupon bonds this way, what does this indicate? It suggests simply that each coupon can be viewed as the payoff of a zero coupon bond. For instance, the second term on the right hand side of Equation (2) is the valuation of a zero coupon bond of face value  $C$ , received at the end of period 2. This analogy is true for all of the cash flows, including the last ( $T$ -period) payment, which is  $(C + FV)$ . Doing this effectively strips the coupons from a standard coupon-bearing instrument, creating a series of zero coupon bonds. The US government does not issue zero coupon bonds. However, investment companies noticed in the early 80's that there was large potential demand for "stripped" treasuries. In 1982, Merrill Lynch and Salomon Bros. created the first widely traded synthetic treasury zeros. These securities are still popular.

Would the price in Equation (2) be the same as the evaluation based on yields as per Equation (9) in Chapter 2? Theoretically, yes. The yield ( $i$ ) used in Chapter 2

turns out to be a weighted average of the T zero-coupon yields. If the bond market is characterized by many actively traded bonds of differing maturities so that a sufficient

set of zero coupon rates can be backed out, then the two equations should generate the same result. In well-developed bond markets (like those in the US), arbitrageurs stand ready to buy and sell different baskets of bills, notes and bonds looking for an arbitrage profit. This activity quickly erases any pricing differences in the two equations. However, in smaller bond markets around the world we see incomplete bond maturity sets and/or trading liquidity problems, which may impede the arbitrage process.

### 3.4 Implied forward rates

Recall that  $R_T$  refers to the equivalent one-period interest rate which would apply over the entire maturity of a particular zero coupon bond. For example,  $R_2$  refers to an average one-period rate applied over two periods. Can we break  $R_2$  down into constituent parts, one for each of the two periods that make up the two-period discount calculation? The answer is yes, as we shall see.

Reconsidering Figure 3.1, imagine we are at time  $t = 0$  and wish to invest \$1 for just one period. You can see that the one-period zero coupon rate  $R_1$  and the spot rate ( $i_1$  or  $f_0$ ) are the same. Your dollar would accumulate to:

$$1 + R_1 = 1 + f_0$$

dollars at the end of period 1.

However, imagine that at  $t = 0$  you had wanted to commit your funds for two periods. There are now two ways of writing the resulting accumulated value, and both must be equal:

$$(1 + R_2)^2 = (1 + f_0) \times (1 + f_1)$$

The left hand side of this equation says that because you are tying up your money for two periods, you must earn the two-period yield  $R_2$ . The right hand side suggests that this can also be regarded as two one-period investments, one after the other. You invest at rate  $f_0 (= i_1)$  over the first period. Then you invest the proceeds  $(1 + f_0)$  over the second period at rate  $f_1$  to get  $(1 + f_0) \times (1 + f_1)$ .

Since  $f_0 = R_1$ , we can back out what the  $f_1$  must be by solving the above equation to get:

$$f_1 = \frac{(1 + R_2)^2}{(1 + R_1)} - 1 \quad (3)$$

The whole process can be extended to time periods further into the future, to recover all the implied forward rates in terms of the zero coupon rates.

In general, if we have the zero coupon rates for time  $t + 1$  and for time  $t$ , we can calculate the forward rate  $t$  periods from now. The implied forward rate,  $f_t$ , for maturity  $t$  in the future is given implicitly by:

$$(1 + R_{t+1})^{t+1} = (1 + R_t)^t \times (1 + f_t)$$

or explicitly by:

$$f_t = \frac{(1 + R_{t+1})^{t+1}}{(1 + R_t)^t} - 1 \quad (3')$$


Hence, if you have a series of zero coupon rates ( $R_1, R_2, R_3 \dots$ ), you can always back out a corresponding series  $f_0, f_1, f_2$  of implied forward rates.

There is an approximate averaging relationship between the zero coupon rates and the implied forwards:

$$R_t \approx \frac{1}{t} \times (f_0 + f_1 + \dots + f_{t-1}) \quad (4)$$

The wavy equal sign means “approximately equal to.” The approximation becomes exact as the compounding interval gets smaller and smaller. At any rate, you can think of the  $R$ 's as the average of the  $f$ 's. An exact (rather than approximate) relationship says that  $(1 + R_t)^t = (1 + f_0)(1 + f_1) \dots (1 + f_{t-1})$ . If you are mathematically inclined, you will know this means that the  $t$ -period zero coupon accumulation factor is the geometric average of the one-period accumulation factors  $1 + f_0, 1 + f_1, \dots (1 + f_{t-1})$ .

The term structure can also be portrayed in terms of the  $f$ 's, that is, by plotting  $f$  against forward time  $t$ . This is called a *forward rate term structure*. In many ways, it is better than the term structure based on the zeros, as it gives you the market's estimates of the future spot rates *at* the particular time interval. In business, many think that marginal rates are always more useful than average rates, whether you are looking at interest rates or tax rates.



**Figure 3.3. Implied forward versus zero coupon term structure**

Figure 3.3 plots a forward rate term structure and the corresponding zero coupon term structure (the curve with a hump) on the same diagram. The averaging relationship referred to earlier means that the graph for  $R$  tends to follow that for  $f$ . This is why the peak in  $R$  occurs after the peak in  $f$ .

### 3.5 Forwards, futures and no-arbitrage

Are forward rates and forward prices related? The following discussion will show us that they are. But before discussing this relationship, we need to review some of the concepts from Chapter 2.

Recall that forward contracts are an agreement calling for future delivery of an asset at a price agreed upon at the inception of the contract. A futures contract is a forward contract that is exchange traded and marked to market, i.e. valued every day. To assist tradability, the futures contract is standardized for units, quality and quantity. As the following example will show, forward rates bear some relationship to forward or futures prices on these instruments.

Recall also the definition of *hedging*—investing in an asset to reduce risk. The concept of a perfect hedge is integral to the discussion that follows. If you are long an asset, the hedge would be to short a futures (forward) contract. If you were

short the asset, the hedge would be to go long on a futures (forward) contract on the asset. If a hedge is perfect, you reduce the price risk of the portfolio to zero.

For the purposes of this example we will discuss contracts on unit discount bonds and ignore the trading complexities of the various interest rate contracts discussed in Chapter 2. By doing this, we ignore the peculiarities of each of these contracts while pointing out the relationship between forward (futures) prices and forward rates.

Imagine a simple two-period scenario. Period one spans the interval between  $t = 0$  and  $t = 1$ , while period two spans the interval from  $t = 1$  to  $t = 2$ . Let's suppose that today ( $t = 0$ ) you enter into a contract to deliver a one-period unit discount bond one period from now (that is,  $t = 1$ ) for an agreed upon price,  $F$ . In other words, after the passage of one period of time, you will deliver to the other person in the contract a one-period unit bond, and the other person will give you  $F$  in exchange. The next topic of our discussion is the *no-arbitrage* price  $F$ , which is the price to be quoted in the market now so that nobody can make money for certain when seen as of time 0, the current time.

First, some notation:

$F$  = price of a one period bond agreed at  $t = 0$  for delivery at  $t = 1$

$P_1$  = price of a one period unit discount bond at  $t = 0$

$\tilde{P}_1$  = price of a one-period unit discount bond at  $t = 1$

$P_2$  = price of a two-period unit discount bond at  $t = 0$

Note the distinction between  $P_1$  and  $\tilde{P}_1$ . The first refers to the price of the one-period bond at the current time ( $t = 0$ ). The second refers to the price of the one-period bond at  $t = 1$ . The tilde symbol over the  $P$  in  $\tilde{P}_1$  indicates that it is a random variable (one whose outcome is yet unknown). Indeed,  $\tilde{P}_1$  is unknown at time  $t = 0$ .

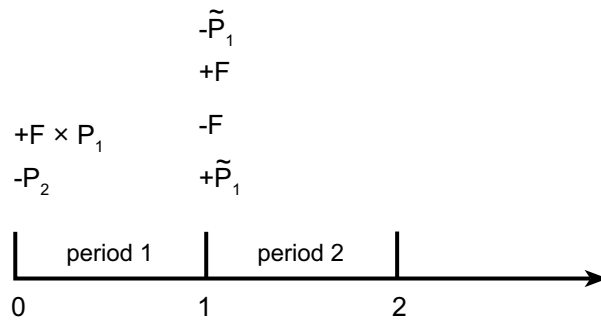
Now back to the problem of determining  $F$ . You have sold a one-period bond forward, so now you have an exposure to  $\tilde{P}_1$ , the price of the one-period bond you will have to use to deliver on your forward commitment. One characteristic of arbitrage is that your position be risk-free, so that at inception, you have to cover



the above end-of-period exposure to  $\tilde{P}_1$ . Another characteristic is that true professional arbitrage plays should cost no money of your own. So how do we manufacture this?

First, we have to cover the exposure described above. To do so, go long (that is, buy) a two-period discount bond now. At the end of the period, it will have become a one-period discount bond. Being long in this bond will cancel exposure to  $\tilde{P}_1$ .

Of course, going long will cost money, as we have to purchase the two-period bond. To fund this purchase, we will borrow against the cash flow of  $\$F$  to be received at time  $t = 1$ . This one-period borrowing is done via the issue of  $F$  one-period unit bonds at price  $P_1$ . (Note that  $P_1$  is the known price at  $t = 0$  of a one-period bond). Thus, we borrow  $\$(F \times P_1)$  which is noted as a cash inflow (+) while the purchase of the two-period unit bond is noted as a cash outflow (-). Let's draw a timeline to illustrate these transactions.



The following table lists the contents of this timeline and shows the detail of the cash flows at  $t = 0$  and at  $t = 1$ .

Transactions at time 0	Cash Flow in \$	
	$t = 0$	$t = 1$
Buy 2 period unit bond	$-P_2$	$+\tilde{P}_1$
Sell F 1 period unit bonds	$+F \times P_1$	$-F$
Enter into Forward Contract	0	$-\tilde{P}_1$
Receive F at $t = 1$	0	$+F$
Total	$+F \times P_1 - P_2$	0

Because of the way this example is constructed, the actions in the present ( $t = 0$ ) result in a net zero cash flow in the future ( $t = 1$ ). What should a portfolio that is worth \$0 in the future be worth today? The answer must be zero. We can take this result and set the total cash flow at  $t = 0$  to \$0, or  $F \times P_1 - P_2 = 0$ . Rearranging the terms, we get:

$$F = \frac{P_2}{P_1} = \frac{\text{price of two period bond}}{\text{price of one period bond}} \quad (5)$$

We know from Equation (1) that:

$$P_1 = \frac{1}{(1 + R_1)}, \text{ and}$$

$$P_2 = \frac{1}{(1 + R_2)^2},$$

so,

$$\frac{P_2}{P_1} = \frac{1 + R_1}{(1 + R_2)^2}.$$

Equation (3) allows us to back out the one-period forward rate one period from now. This, along with the equation above, tell us that the forward or futures price,  $F$ , is equal to:

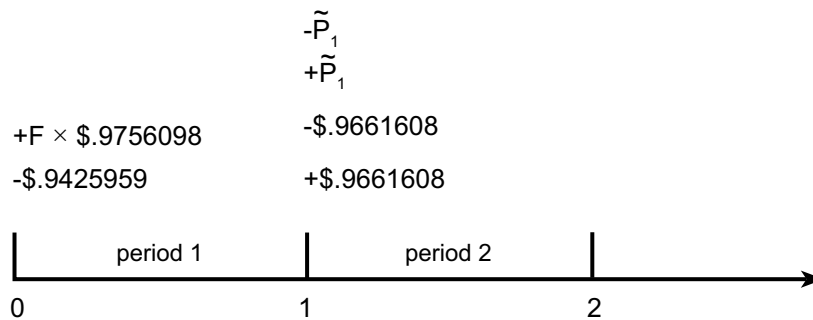
$$F = \frac{1}{1 + f_1} \quad (6)$$

This is why  $f_1$  is called the implied forward rate. It represents a market certainty equivalent discount rate formed at  $t = 0$ , to enable it to price now a one-period discount bond for delivery at instant  $t = 1$ , which is future time. More generally, the forward rates  $f_0, f_1, f_2$ , etc., represent a pricing device by which the forward component of any current price can be calculated.

Returning to the example above, let's insert some actual numbers.

$$\begin{aligned}
 R_1 &= 2.5\%; & P_1 &= \$0.9756098 \\
 R_2 &= 3.0\%; & P_2 &= \$0.9425959 \\
 F &= 0.9661608
 \end{aligned}$$

The two cash transactions at  $t = 0$  are to buy a two-period discount bond now at  $\$0.9425959$ , and issue  $F$  one-period unit bonds at  $\$0.9756098$ . Thus, we have borrowed  $\$(F \times \$0.9756098)$  and this is noted as a cash inflow (+), while the purchase of the two-period unit bond is noted as a cash outflow (-). The value of  $F$  we have calculated as  $0.9661608$ , while that of  $\tilde{P}_1$  is still unknown, so we simply use a symbol. We will redraw the timeline to incorporate these cash flows.



Transactions at time 0	Cash Flow in \$	
	$t = 0$	$t = 1$
Buy 2 period unit bond	$-\$.9425959$	$+\tilde{P}_1$
Sell $F$ 1 period unit bond	$+F \times \$.9756098$	$-F$
Enter into Forward Contract	0	$-\tilde{P}_1$
Receive $F$ at $t = 1$	0	$+F$
Total	$+F \times \$.9756098 - \$.9425959$	0

The no arbitrage condition tells us that  $F \times \$0.9756098 - \$0.9425959 = 0$ . Rearranging terms we see that:

$$F = \frac{P_2}{P_1} = \frac{\$.9425959}{\$.9756098} = .9661608$$

Using Equation (6) and  $F = .9661608$ , we find that  $f_1$  is 3.50244%. Of course, our current purpose is illustrative, and we do not need to go through all these calculations in order to obtain the implied forward rates! To find the rates, simply use Equation (3) or (3r) above.

### 3.6 Computing zeros and forwards

The correct way to compute the zero rates is from an actual series of zeros or from bonds that have been stripped and repackaged as zeros. To estimate the entire yield curve and forward rates we would need an entire series of such bonds. What if there are not enough zeros or strips? This is not a problem if you have prices on a set of coupon bonds that span the maturity horizon for the yield curve (as in the US—there are over 150 government bonds listed in the *Wall Street Journal* with maturities ranging from one day to 30 years). At any rate, suppose you have three coupon bonds with coupons  $C_1, C_2, C_3$  and maturities one, two and three periods. The market quotes them at prices  $P_1, P_2, P_3$ , respectively. In terms of the zeros, the bond prices can be decomposed and written (with subscripts denoting the bond):

$$P_1 = \frac{C_1 + FV}{1 + R_1} ;$$

$$P_2 = \frac{C_2}{1 + R_1} + \frac{C_2 + FV}{(1 + R_2)^2} ;$$

$$P_3 = \frac{C_3}{1 + R_1} + \frac{C_3}{(1 + R_2)^2} + \frac{C_3 + FV}{(1 + R_3)^3} .$$

Notice how the bonds are valued. As observed in Section 3.3, we are decomposing them into their constituent cash flows. A coupon received at time three is itself just like a zero coupon bond of face value  $C_3$  maturing at time three.

Once the market quotes the three prices, you can solve for the three zero coupon rates implied by those prices. Just solve the three simultaneous equations above. Of course, once you have the zeros, getting the implied forwards is straightforward.

Instead of using prices, market yields are often used to back out zeros. We have used yields – or internal rates of return – very extensively in Chapter 2. Up to this

point, however, we have not mentioned the relationship between these yields and the zeros or implied forwards, so we need a digression on this.

### Yields, zeros and swaps

You can also do a term structure based on yields. Suppose you had three coupon bonds, as above, being quoted at market prices  $P_1$ ,  $P_2$ , and  $P_3$ . You know the yields on these bonds will all be different. We will use the symbol  $\rho$  for the yield, as it is often used for the internal rate of return, and the yield is just an IRR. So the three yields  $\rho_1$ ,  $\rho_2$ ,  $\rho_3$  are defined implicitly by the equations (the subscripts denote the particular bond):

$$P_1 = \frac{C_1 + FV}{1 + \rho_1} ;$$

$$P_2 = \frac{C_2}{1 + \rho_2} + \frac{C_2 + FV}{(1 + \rho_2)^2} ;$$

$$P_3 = \frac{C_3}{1 + \rho_3} + \frac{C_3}{(1 + \rho_3)^2} + \frac{C_3 + FV}{(1 + \rho_3)^3} .$$

In other words, once you know the prices, you can recover the three yields to maturity, namely  $\rho_1$ ,  $\rho_2$ , and  $\rho_3$ , for the different maturities. You can also plot these yields against time, which would give you a *yield-based term structure*—yet another version of the term structure.

We have seen that one can switch readily between zeros and implied forwards – they are simply different aspects of the same thing. Can you also do the same as between yields and zeros? The answer is “yes and no.” It depends upon whether your coupon bonds, as in the example above, are trading at par. You will recall that a bond trades at par if the current price,  $P$ , is equal to the face value,  $FV$ . It turns out that if the bonds are not trading at par, then there is no simple relationship between the yields and the zeros.

Indeed, the yield to maturity depends on the coupon to face value ratio (the coupon yield), so the maturity alone is not sufficient information to determine the yield. The yield depends on various characteristics of the particular bond. This makes matters a bit inconvenient. If you have two three-year bonds with different coupons, they will have quite different zero coupon-based prices, and possibly different yields as well. There is no simple relationship, independent of the coupon, between yields and zeros. In fact, if you do price different three-year bonds using the same the

three-year yield computed off just one of them, you could face arbitrage by someone who has calculated the price correctly using zeros. Recall that zeros are independent of coupons, face values, etc., and are the same for every bond. In this sense, the zeros—not the yields to maturity—would be the correct way to price the bond.

If you calculated the yield to maturity on a bond that happened to be trading at par, you could in fact recover the zeros. One example where trading at par for a fixed interest instrument always occurs is with an interest rate swap, just placed on the books. An interest rate swap is an agreement between two parties, whereby each contracts to make interest payments to the other on specific dates in the future.

The standard interest rate swap involves an exchange of fixed rate payments, which we will refer to as *coupons*, determined at the inception of the swap, for a set of floating rate payments determined on specified rate setting dates. Both fixed and floating payments are calculated on a notional principal amount that never changes hands. These floating rate payments are typically based on the six-month LIBOR rate (the rate for six-month interbank dollar deposits set in London) on each rate set date. In fact, swaps can come in all sorts of flavors, but this is a basic or “plain vanilla” swap.

A swap has a zero value when placed on the books. It represents a free exchange of one sort of payment for another. The fixed side (at inception) is valued as though it were a coupon bond selling at par. The floating side represents market indifference between such a coupon bond and the series of future floating rate payments. So here it is: *the yield on the fixed side of the swap at its inception can be taken as a par yield.*

Swaps are quoted in standard maturities of one, two, three, four, five, seven and ten years, with in-between maturities usually calculated by applying a linear interpolation. For example, a six-year swap would typically be quoted with a fixed rate half way between the five- and seven-year par yields. Using the yield to maturity of the many different swap maturities, you can calculate backwards to find the zero coupon rates. In the next section, we give you an algorithm, suitable for use in a spreadsheet, to do just this.

### 3.7 Algorithm: computing zeros and forwards from swap data

“Plain vanilla” (or standard) swaps are quoted on the basis of a constantly changing coupon and equivalent yield, so they are traded at par when transacted. They are, however, an over-the-counter product, and even swaps quoted on the same day and same maturity are not usually homogeneous. As market prices of bonds, futures, and so on move from one minute to the next, so does the swap rate (which represents both the coupon and the yield). As we’ve discussed, you cannot expect that a set of cash flows with the same maturity but different coupons can be correctly priced at the same yield. Sometimes this will turn out to be the case, but more often it will not. Hence, it is impossible to value one-day-old (even one-minute-old) swaps using yields to maturity.

Swaps are always valued by applying the current swap zero-coupon curve. At the outset, the zero-coupon valuation gives the same answer as a valuation using the yield, but not after time ticks onward and swap rates move. At any given point in time, new swaps are quoted on a yield basis, often as a spread over a government bond yield curve, and the current swap zero curve can be obtained from the swap “yields” by the following iterative process.

#### Backing out the swap zeros

Suppose that we find that the US swap market is currently quoting the following swap rates, displayed in Table 3.1, on a percentage per annum basis. We have just listed mid-rates out to three years to simplify matters. (The full range can go out ten years or more for top credit counterparties.) The following algorithm also abstracts from the complexities of day-count conventions by using a semi-annual coupon, equal to half the annual coupon rate times the notional principal (bond basis). The semi-annual yields are just the current par swap rates divided by two.

In the US market, swaps are frequently quoted with an annual Actual/360 coupon, with the floating leg based on six month LIBOR, paid semi-annually. Our simplified example illustrates the methodology for calculating the zero-rates without getting bogged down in day-count issues. It can readily be adjusted to handle any of the various day-count conventions used in the different swap markets.

**Table 3.1** Par swap rates

Maturity (years)	Period (t)	Par Swap Rate (% p.a.)
0.5	1.0	9.53 (LIBOR rate)
1	2.0	9.08
1.5	3.0	8.84
2	4.0	8.6
2.5	5.0	8.5
3	6.0	8.4

**Step 1** The first point on the zero curve is the easiest. It is simply the six-month bank bill rate which is itself a zero coupon rate. The first discount factor  $d_1$  is therefore:

$$d_1 = \frac{1}{1 + \frac{9.53}{200}} = 0.95451725.$$

**Step 2** If we take the fixed side of a current one-year semi-annual swap with face value of 100, and use the discount factor  $d_1$  to value its first cash flow, we can solve for the second zero rate. It must be the rate which values the second cash flow at an amount that will bring the total value of the fixed cash flows (including the final flow of principal) back to par—i.e., equal to 100.

That is:

$$100 = \frac{C_1}{1 + \frac{9.53}{200}} + \frac{C_1 + 100}{\left(1 + \frac{R_2}{200}\right)^2}.$$

Here  $C_1$  is the “coupon” for a current one-year swap, which per hundred face value on a swap paying on a semi-annual basis is:



$$C_1 = \frac{9.08}{2} = 4.54 \quad .$$

We can plug the value for  $C_1$  into (8) as follows:

$$100 = \frac{4.54}{1 + \frac{9.53}{200}} + \frac{4.54 + 100}{\left(1 + \frac{R_2}{200}\right)^2} \quad ,$$

and then rearrange to get:

$$\frac{1}{\left(1 + \frac{R_2}{200}\right)^2} = \left(100 - \frac{4.54}{1 + \frac{9.53}{200}}\right) \times \frac{1}{104.54} \quad .$$

Calculating the right side of the equation, we obtain:

$$d_2 = \left(\frac{1}{1 + \frac{R_2}{200}}\right)^2 = 0.91511854 \quad .$$

This is our second discount factor.

The discount factor is very useful as it stands for discounting one-year cash flows, but we will take the next step and solve for  $R_2$  as well:

$$R_2 = \left(\frac{1}{.91511854^{0.5}} - 1\right) \times 200 = 9.069808 \quad .$$

**Step 3** As the 1.5 year swap rate is 8.84% p.a., the semi-annual coupon will be:

$$C_{1.5} = \frac{8.84}{2} = 4.42$$

We proceed now in a similar manner to that of Step 2 to find the third discount factor  $d_3$ , and then the third zero rate,  $R_3$ .

Summarizing in Table 3.2, our zero curve at this point:

**Table 3.2 Swap Zero curve, Step 3**

Period	Par rate	Zero rate	Discount factor
1.0	9.53	9.53	.95451725
2.0	9.08	9.0698	.91511854
3.0	8.84	$R_3$	$1/(1+R_3/200)^3$

The fixed side cash flows of our current 1.5-year swap are valued as:

$$100 = 4.42 d_1 + 4.42 d_2 + \frac{4.42 + 100}{\left(1 + \frac{R_3}{200}\right)} .$$

We can again solve for the third zero rate by starting with the discount factor:

$$d_3 = \frac{1}{\left(1 + \frac{R_3}{200}\right)^3} = (100 - 4.42(d_1 + d_2)) \times \frac{1}{104.42} .$$

Inserting  $d_1$  and  $d_2$  which we've already computed, we get:

$$d_3 = \frac{1}{\left(1 + \frac{R_3}{200}\right)^3} = 0.8785310 .$$

In turn, we obtain:

$$R_3 = \left( \frac{1}{.8785310^{0.3333}} - 1 \right) \times 200 = 8.8226618 \quad .$$

**Step 4** Next, find the fourth discount factor and zero rate, using the three preceding discount factors and applying the same process as above. The two-year (fourth period) swap rate is 8.60% p.a., giving a coupon rate of:

$$C_{2.0} = \frac{8.60}{2} = 4.30 \quad .$$

**Table 3.3 Swap Zero curve, Step 4**

Period	Par rate	Zero rate	Discount factor
1.0	9.53	9.53	.95451725
2.0	9.08	9.0698	.91511854
3.0	8.84	8.82266	.8785310
4.0	8.60	$R_4$	$1/(1+R_4/200)^4$

We can express the value for the fixed cash flows of a current two-year swap as:

$$100 = 4.30 d_1 + 4.30 d_2 + 4.30 d_3 + \frac{4.30 + 100}{\left(1 + \frac{R_4}{200}\right)^4} \quad ,$$

and find the two-year (fourth period) discount factor:

$$d_4 = \frac{1}{\left(1 + \frac{R_4}{200}\right)^4} = (100 - 4.30 (d_1 + d_2 + d_3)) \frac{1}{104.30} = .8454735 \quad .$$

and then find the fourth zero rate:

$$R_4 = \left( \frac{1}{.8454735^{0.25}} - 1 \right) \times 200 = 8.571517 \ .$$

You can see how it goes: you arrive at the discount factors  $d_1, d_2, d_3, \dots$  recursively from the previous ones. Once you have each  $d_t$ , you can easily obtain the associated zero curve rate  $R_t$ . Now you can test your understanding of these concepts by recovering the remaining  $d_t$  and  $R_t$  for the data of Table 3.1. You should be able to calculate the discount factors and zero rates as shown in Table 3.4.

**Table 3.4 Swap Zero curve to three years**

Period (t)	Par rate	Zero rate $R_t$ %	Discount factor
1.0	9.53	9.53%	.95451725
2.0	9.08	9.06981	.91511854
3.0	8.84	8.82266	.87853103
4.0	8.60	8.57152	.84547347
5.0	8.50	8.46868	.81272929
6.0	8.40	8.36320	.78208491

### Finding the forward rates

For the last part of this example, we will calculate the one-period forward curve,  $f_0, f_1, \dots, f_n$  from the zero rates above.

Retrieving Equation (3) from Section 3.4 and modifying it slightly to adjust for our semi-annual data, we can apply it to the zero rates we have calculated.

$$f_1 = \frac{(1 + R_2)^2}{(1 + R_1)} - 1 \quad . \quad (7)$$

Recall that the first forward rate,  $f_0$ , is simply the zero rate,  $R_1$ , which is 9.53% on Table 3.4. The next forward rate,  $f_1$  is found by plugging  $R_1$  and  $R_2$  from the same table into Equation (7a) as follows.

$$f_1 = \left( \frac{\left(1 + \frac{R_2}{2}\right)^2}{1 + \frac{R_1}{2}} - 1 \right) \times 2 \quad . \quad (7a)$$

$$f_1 = \left( \frac{\left(1 + \frac{0.090698}{2}\right)^2}{1 + \frac{0.093}{2}} - 1 \right) \times 2 = 0.0861063 \quad \text{i.e. } 8.6106\% \quad .$$

To find the remaining one-period forward rates, we need a more general form of (7a), which we will call (7b):

$$f_{t-1} = \left( \frac{\left(1 + \frac{R_t}{2}\right)^t}{\left(1 + \frac{R_{t-1}}{2}\right)^{t-1}} - 1 \right) \times 2 \quad . \quad (7b)$$

The third forward rate is then:

$$f_2 = \left( \frac{\left(1 + \frac{0.0882266}{2}\right)^3}{\left(1 + \frac{0.090698}{2}\right)^2} - 1 \right) \times 2 = 0.083292 \quad \text{i.e., } 8.3292\% \quad .$$

The fourth forward rate,  $f_3$  is:

$$f_3 = \left( \frac{\left(1 + \frac{0.085715}{2}\right)^4}{\left(1 + \frac{0.0882266}{2}\right)^3} - 1 \right) \times 2 = 0.078199 \quad \text{i.e. } 7.8199\% .$$

The one-period swap forward rates and zeros derived from the original (three year) swap par curve of Table 3.1 are collected in Table 3.5 below. You can test your understanding of this section by attempting to produce the fifth- and sixth-period rates yourself. Better yet, produce the curves out to five years using the augmented data given in Table 3.6 in the exercises at the end of this chapter. Such a set of calculations is most easily performed using a spreadsheet program such as Excel or Lotus. If you are not familiar with these programs, particularly with respect to financial calculations, you may find it helpful to skip ahead to Chapter 5, which provides an introduction to spreadsheets.

**Table 3.5 Swap Zero and Forward curve to three years**

Period (t)	Par rate %	Zero rate $R_t$ %	Forward rate $f_{t-1}$ %
1.0	9.53	9.53%	9.53%
2.0	9.08	9.06981	8.61063
3.0	8.84	8.82266	8.32925
4.0	8.60	8.57152	7.81989
5.0	8.50	8.46868	8.05783
6.0	8.40	8.36320	7.83659

### 3.8 Concluding remarks

Modern fixed-interest pricing theory operates using zeros, and especially, implied forward rates. A popular approach to modeling the term structure is by assuming that there are only a few primary drivers in determining forward rates. These factors are assumed to drive not only the forward rates, but flow through these to the entire term structure and the pricing of fixed interest products.

Although we have presented our discussion in terms of discrete time, these concepts are more naturally cast in terms of continuous time discounting. This notion was introduced in Section 2.5, but there the term structure was effectively assumed to be flat, so there was only one interest rate to worry about. Where there are several forward rates, the representative discount factor – equivalent to the price of bond of continuous maturity  $T$  – is given by:

$$P_T = e^{-\int_0^T f_t dt} .$$

In this equation, the discounting process occurs over a sequence of small intervals from  $t$  to  $t + dt$ , and at each interval, the instantaneous forward rate  $f_t$  is applied in forward time  $t$ . The zero coupon rate for the interval  $(0, T)$  is defined by:

$$R_T = \left( \frac{1}{T} \right) \int_0^T f_t dt .$$

Bearing in mind that the integral is just a sum, you can see that the zero rate is just the arithmetic mean of the instantaneous forward rates. We could also write the discounting formula as:

$$P_T = e^{-T R_T} ,$$

and the relevance to the discussion of continuous discounting in Chapter 2, Section 5, should be clear.

---

We have not discussed the topic of the term structure in much detail, nor explained what determines the way the term structure curve shifts through time (as it has certainly done recently in many countries). There are several theories about the term structure. One long-standing theory is that the forward rates can be regarded as *expectations* of future spot rates, conditional on information available at time zero. For instance, if you expect that two years from now, the Federal Reserve will be adopting a tighter monetary stance which would affect short-term interest rates, so you are expecting the spot rate in two years time to be high. Under the pure expectations theory, you revise your current forward rate upwards for two years out. The theory is sometimes augmented to include risk influences, which generally increase with forward time, since future events are viewed as being more risky.

We have not yet finished with the topic of interest rates, which will become important as we go on to discuss foreign exchange. However, what has been covered so far will serve as an introduction to this huge and multi-faceted topic, which is of immense importance in financial management.



The basic data needed for the following questions is contained in the table below, an augmented version of Table 3.1 which appeared earlier in the chapter.

**Table 3.6 Par swap rates**

Maturity (years)	Par Swap Rate (% p.a.)
0.5	9.53 (6-month LIBOR rate )
1.0	9.08
1.5	8.84
2.0	8.60
2.5	8.50
3.0	8.40
3.5	8.35
4.0	8.30
4.5	8.27
5.0	8.24

1. Use the swap zero coupon rates to value two coupon bonds, each with face value \$1,000,000 with maturity five years but with differing six-month coupon rates as follows:
  - (i) A coupon rate of 10%
  - (ii) A coupon rate of 5%

(Begin by computing the zero rates for the last four time periods.)

2. Given your prices from Q1, calculate the yields to maturity for each bond. Are they the same?
3. In Chapter 2, we determined the duration of a set of cash flows by using the yield as the valuation tool. You can also find the duration by using the zero coupon rates; the result is called *Fisher-Weil duration*. In this method, simply value each cash flow by using the zero coupon rate for that particular cash flow bucket (that is, if the cash flow occurs at bucket  $t$ , use  $R_t$ ). Likewise, the overall value of the cash flows, or instrument, if they take this form has to be established by using the zero coupon rates.

Work out the Fisher-Weil durations for each of the two coupon bonds of Q1, using the zero coupon rate data.

*Note:* Ordinary, or *Macaulay*, duration is the change, or proportional change, in the value of the portfolio or instrument following a change in the yield. Fisher-Weil duration can be interpreted as the change in the value of the portfolio following a parallel change in the term structure of the implied forward rate—i.e. imagining that each forward rate shifts upwards or downwards by exactly the same amount.

4. Let's do a bit of swap valuation. Well known lunchtime bon viveur Sherry van Plonk puts a five-year swap on her books at the set of zeros implied by the data in Table 3.6. She is receiving the fixed side. Then she goes off to lunch, but does not bother to hedge the exposure under the "she'll be right, mate" precept. Three hours later she returns, just in time to hear that Alan Greenspan has had a bad lunch, and the zeros term structure has changed as follows:

New rates (still applying for six month intervals):

$$\begin{aligned}R_1 & 10.20\% \text{ p.a.} \\R_2 & 9.55 \\R_3 & 9.26 \\R_4 & 9.01\end{aligned}$$

and the remaining rates as before lunch.

How much has Ms. van Plonk won or lost on her swap?

5. During their very enjoyable lunch, well known tax consultant Darth Evader has suggested to Ms. van Plonk that she do him a swap whereby he will make a single payment up front, in return for a set of floating payments out to five years. The nominal or face value of the swap will be \$1m.
- (i) Using zeros for the valuation, what should Ms van Plonk be asking him to pay?
- Assume pre-lunch interest rates and ignore commission and other transaction costs.
- (ii) Would such a swap be more or less exposed to the post-lunch change in interest rates?
6. Find the Fisher-Weil duration of the swap in Question 4 above (valued using the post-lunch rates).



---

# Four

## FX Spot and Forwards

Around 1973, the world's financial markets moved to a system of floating exchange rates, which means that the price of one currency is free to move around relative to that of other currencies. Even the *Hitchhiker's Guide to the Galaxy* chats about the exchange rate between the Altairian Pobble Bead and the Triganic Pu— this is an intergalactic phenomenon! Knowing how exchange rates are quoted is important for any manager who deals internationally—which today means just about every manager.

In addition to looking at spot prices, we will also discuss trading and pricing of FX forwards, which are contracts to buy or sell foreign exchange at some designated time in the future. These FX contracts are a basic part of the risk management of foreign exchange transactions (such as accounts payable or receivable) where the change in value of the transaction currencies will affect your bottom line. You will see that forward contracts can mitigate the effect of these price swings. We will show you that the price function of a forward is multidimensional—based on interest rates inside countries, as well as the corresponding spot exchange rates between them.

#### 4.1 Notation and definitions

Ask	The price at which a dealer or market maker is willing to sell. It is the price at which everyone else can buy. Sometimes called the offer price.
Bid	The price at which a dealer or market maker is willing to buy—i.e. the price at which the customer is able to sell.
Commodity currency	The currency being normalized to one unit in a quotation; more simply, the currency being purchased (or sold) in terms of the other currency.
Cross rate	Rate of exchange between two currencies that are less actively traded. Often applied to currencies that are traded against one another, neither of which is the US dollar. These rates need to be calculated from exchange rates quoted against the US dollar.
Exchange rate	The exchange rate is the price of one currency in terms of another.
Forward rates	The exchange rate for currency transactions for future value that take place normally at least three business days ahead. (Can be two business days ahead for some North American currencies, such as the Canadian dollar.)
Inversions	The algebraic inversion of a given exchange rate. Bid prices invert to ask prices and ask prices invert to bid prices.
Mid-rate	The average of the bid and ask prices.
Numeraire	Another name for the commodity currency.
Spot rates	The exchange rate for currency transactions for current value with settlement taking place in two business days. For some transactions in North America, settlement time can be one day..(e.g. CDN)
Terms currency	In a quotation, the currency used to buy or sell against the commodity currency.

## 4.2 Spot exchange rate quotations

The exchange rate is the rate at which one currency is exchanged for another. Since we have one author from the US and another in New Zealand, we will start our discussion with the US dollar (USD) vis-à-vis the New Zealand dollar (NZD). Recently we could exchange one New Zealand dollar for  
 NZD1 = USD 0.4290

As in any pricing or exchange scheme, there are two things to be considered. In the case of exchange rates, both of the items are money, so we need to standardize on one of the two. Above, the NZD is the *numeraire* or *commodity unit of measure*, which is quoted in terms of USD. The USD is called the *terms currency*. You could think of the NZD as though it were a pound of kiwifruit being traded for a certain number of U.S. dollars (or cents). Hence, the NZD is the commodity currency in this form of quotation.

The exchange rate can be expressed as:

$$e = \frac{Q_{\text{terms}}}{Q_{\text{commodity}}},$$

where  $Q_{\text{terms}}$  is the amount of the terms currency,  $Q_{\text{commodity}}$  is the amount of the commodity currency and  $e$  is the exchange rate, the amount of terms currency per unit of the commodity currency.

The commodity currency is best thought of as the object being bought or sold, using the terms currency. Profit or loss on currency trading therefore naturally arises in the *terms* currency. If the price of the kiwi goes up to USD 0.4290 from USD 0.4270, we instinctively know in the U.S. that the price of this commodity went up. That is the normal way of thinking about it. If, however, someone were to quote it the other way – as NZD 2.3310 from NZD 2.3419 – we wouldn't intuitively know whether the price went up or down. If we always use the domestic currency (in our case the USD) as the terms currency, the exchange seems quite clear. In the US, we usually do look at the price of other currencies this way, but once you leave our shores, this is not always the case. We'll explore this issue below.

You can test your new-found understanding in the context of another common practice outside of the US. Often exchange rates are written as—say—USD/JPY, to use the Japanese Yen as an example. This means that 1 USD equals a certain number of JPY, making the USD the commodity currency in this case. On

most quote screens, the first currency listed is the commodity currency, and the second is the terms currency.

### Standard short forms

When you see Foreign Exchange quoted in newspapers in the United States, you usually see the name of the country with the name of the currency in parenthesis. See Table 4.3 in Section 4.5 for an example. This is not the case in most other countries. When you look at newspapers, FX screens, or notice boards in other countries, you'll find the various currencies are typically reported in the standard short forms, or derivations, as illustrated in Table 4.1 below. These abbreviations are based on the Reuters classifications.

**Table 4.1 Foreign exchange conventions: short forms**

Short Form	Country (Region)	Short Form	Country (Region)
ARS	Argentina	GBP	UK
AUD	Australia	HKD	Hong Kong
CAD	Canada	JPY	Japan
CHF	Switzerland	NZD	New Zealand
CNY	China	SGD	Singapore
EUR	Germany	USD	United States
EUR	European	ZAR	South Africa

### FX Conventions: direct versus indirect quotes

#### *Direct quotes:*

In dealings with non-bank customers, banks in most countries use the USD as the commodity currency.

#### *Indirect quotes:*

Banks of former British Empire countries (e.g. AUD, NZD, GBP) generally quote currencies using the domestic currency as the commodity and the USD as the terms currency.

**Table 4.2 FX conventions: quotations**

Commodity Currency	Terms Currency
NZD1	= USD 0.4484
AUD1	= USD 0.5392
USD1	= JPY 126.86
USD1	= CHF 1.5760



*Exceptions to this rule:*

For transactions in the US (and for the Canadian Dollar), US banks use the USD as the terms currency. However, US banks follow world conventions for transactions outside of the US.

*US Newspaper quotes:*

Quotes in all US newspapers give the price of non-US currencies in terms of the US dollar—a price that all Americans are used to. Seldom do you see the Reuters convention listed. The mid-price is usually the only price given, along with the country name and name of the currency. Some papers like the *Wall Street Journal* will also give the mid-price inversion (see below).

**Significant digits**

The number of significant digits used in foreign exchange calculations is also governed by convention. For rates of exchange close to unity it is conventional to express the rate of exchange to four decimal places. Other rates of exchange have their own conventions; for instance, USD/JPY is quoted to two decimal places.

**Bid and offer rates of exchange**

A market maker in foreign exchange quotes two prices—the bid and the ask. For example, the NZD/USD rate quoted may be 0.4280-0.4290. The first price (0.4280) is the price at which the trader will buy the commodity currency, which is NZD. As we discussed above, this is the bid price. The second price (0.4290) is the price at which the trader will sell the commodity currency—the ask or offer price. Note that the smaller figure always comes first and the larger figure second. This is a universal convention, which makes exchange rate listings easier to interpret, as we will see shortly.

An axiom for making money is buying at a lower price than you sell. Market makers (dealers) intend to make a profit by buying low and selling high. The *spread*, or difference between a dealer's bid and ask, gives some protection against adverse exchange rate movements, and augments any favorable ones. Part of the business of being a trader is the requirement to make prices on both sides, bid and offer, and traders have to take temporary positions, which may leave them net long or net short in a particular currency for a period of time. The spread between the prices provides them some compensation for the risk taken.

If you split the difference between the bid and the ask, the price is called the *mid-rate*. With the numbers above, the mid rate for the NZD/USD is 0.4285. The mid-

rate is often used in papers such as the *Wall Street Journal*, where the focus is simply on recording how the exchange rate moved (in response to economic news, for example), rather than on actual trading of foreign currency.


In addition to the market quoted rate, economists use a number of constructed rates, such as the effective or real exchange rates. These are index number constructions, typically taking the form of a weighted average of the percentage price movements of a number of currencies, re-expressed in the form of an index. We will not consider these constructed exchange rates much further, but most advanced international finance books cover this issue.

### 4.3 Inversions

Suppose  $\text{NZD/USD} = 0.4280$ . What would  $\text{USD/NZD}$  be? In other words, suppose the USD is now the commodity currency, or that we are adopting the direct basis for the NZD. Is the answer simply  $1/.4280 = 2.3364$ ? If the 0.4280 is a mid-rate, it would be, but if it is a bid or an ask price, we need to be careful.

If the quote is 0.4280-0.4290, the 0.4280 refers to the bid price, the amount the dealer will pay (in USD) for the commodity (the kiwi in our analogy). This is the amount that you get if you are selling New Zealand dollars. Now suppose we reverse the quotation basis to  $\text{USD/NZD}$ . The first figure of the bid/ask pair refers to the situation where the USD is the commodity, so that you are trying to sell USD and buy (not sell) the kiwi. Remember that the bid and ask prices are different.

Therefore, in order to make things correspond, we must use as a starting point the second figure in the  $\text{NZD/USD}$  based quote—this is the one that refers to selling the kiwi. Hence, the figure 0.4290 should be used, making the correct bid for  $\text{USD/NZD}$   $1/0.4290 = 2.3310$ . Figure 4.1 below shows the conversion process:



**Figure 4.1 Inversions**

#### 4.4 Cross rates

In order to arrive at a rate of exchange between two currencies that are not very actively traded, it is necessary to cross the currencies against a common currency. By market convention, the common currency is the USD. It turns out that this crossing process guards against arbitrages even where the two currencies concerned are actively traded. The term *cross rate* is used to refer to quotations of two currencies, neither of which is the USD.

Two main steps in the calculation are:

- 1) Identify the commodity currency and terms currency.
- 2) Identify how the two rates are to be combined.

The Chain Rule provides a (relatively) foolproof method. We'll start by ignoring two-way rates and introduce them later.

#### Example 4.1

Assume that we are given the exchange rate between the New Zealand dollar and the US dollar as: NZD1 = USD 0.4285, while we are told the exchange rate between the US dollar and the Japanese Yen is: USD1 = JPY 127.10. Suppose further that we wish to express NZD as the commodity currency, and the JPY as the terms currency (i.e., NZD1 = JPY??). Now we can use the chain rule.

$$\begin{array}{lcl}
 1) & \text{JPY??} & = \text{NZD1} \\
 2) & \text{NZD1} & = \text{USD 0.4285} \\
 3) & \text{USD1} & = \text{JPY 127.10} \\
 4) & \text{NZD/JPY} & = \frac{(0.4285 \times 127.10)}{1 \times 1} \\
 & & = 54.46
 \end{array}$$

**Figure 4.2 Chain Rule Example**

**Chain rule**

Line 1: Write the equation (backwards):

$$\mathbf{JPY?? = NZD1}$$

Line 2: Write the exchange rate involving the commodity currency and the common currency:

$$\mathbf{NZD1 = USD 0.4285 \text{ (common currency – RHS)}}$$

Line 3: Write the second constituent exchange rate:

$$\mathbf{USD1 = JPY 127.10 \text{ (common currency – LHS)}}$$

Line 4: Multiply the amounts on each side of the equal sign in lines 2 and 3, and divide by the product of the left hand side of the equation.

$$\mathbf{NZD1= JPY 54.46}$$

To see what we are doing algebraically, let's return to the definition of the exchange rate,  $e$ . The definition of  $e$  is  $\frac{Q_{terms}}{Q_{commodity}}$ , so the exchange rates for the currencies above are:

$$e_1 = \frac{USD 0.4285}{NZD1}$$

$$e_2 = \frac{JPY 127.10}{USD 1}$$

In this case, to solve for the cross rate, we need to multiply the two exchange rates to get the correct cross rate. This is not always the case; sometimes you will need to invert one of the exchange rates to get the desired result.

$$e_3 = \frac{JPY ?}{NZD1} = e_1 \times e_2 = \frac{USD 0.4285}{NZD1} \times \frac{JPY 127.10}{USD 1} = 54.46$$

### Two-way cross rates

When we introduce bid and ask prices into the cross rate calculation, we must take care to select the appropriate rates, bearing in mind first of all that the customer always comes out on the losing end of this proposition. If we are not the market makers for currencies, we must always buy high and sell low; that is, we pay the market makers for currency conversion. If, on the other hand, we are calculating the cross rate from the viewpoint of the market makers, we will come out on top in this transaction.

Secondly, recall that if we invert an ask rate, it becomes a bid rate, and if we invert a bid rate, it becomes an ask rate. Finally, in selecting the rates, remember that if we are seeking the ask cross rate, we need to multiply two ask rates together. If we are in search of the bid, we need to multiply two bid rates together to get a bid cross rate. In other words, don't mix and match.

#### Example 4.2

(a) Find the bid rate  $\text{NZD}1 = \text{JPY}???$ , that is, the rate at which a bank in the market will quote to buy NZD and sell JPY. Here are the steps:

(i) Identify the correct rates to be employed.

NZD/USD is quoted at 0.4280-90 (the normal abbreviated form for expressing 0.4280-0.4290)

USD/JPY is quoted at 127.05-15 (meaning 127.05 – 127.15)

Because the NZD is the commodity currency, the bid is the rate at which the quoting bank is prepared to buy NZD and sell USD in the Kiwi market, and to buy USD and sell JPY in the dollar/yen market. In this case, the quoting bank needs to buy the commodity currency in both markets. Hence, in constructing the bid NZD/JPY, we must use bid rates for both NZD/USD and USD/JPY.

(ii) Apply the chain rule.

**First the bank is looking to buy NZD and sell JPY.**

$$\begin{aligned}
 \text{JPY??} &= \text{NZD1} \\
 \text{NZD1} &= \text{USD } 0.4280 \quad (\text{BID SIDE}) \\
 \text{USD1} &= \text{JPY } 127.05 \quad (\text{BID SIDE}) \\
 \text{NZD/JPY} &= \frac{(0.4280 * 127.05)}{(1 * 1)} \\
 &= 54.38 \text{ BID}
 \end{aligned}$$

Algebraically, using the definition for the exchange rate, we calculate:

$$\text{Bid } e_1 = \frac{\text{USD } 0.4280}{\text{NZD } 1}$$

$$\text{Bid } e_2 = \frac{\text{JPY } 127.05}{\text{USD } 1}$$

$$\text{Bid } e_3 = \frac{\text{JPY ?}}{\text{NZD1}} = e_1 \times e_2 = \frac{\text{USD } 0.4280}{\text{NZD1}} \times \frac{\text{JPY } 127.05}{\text{USD1}} = 54.38$$

- (b) Find the ask rate NZD1 = JPY???, that is, the rate at which a bank in the market will quote to sell NZD and buy JPY. This time the quoting bank needs to sell the commodity currency in both markets, so the ask rates are used.

**The bank wishes to sell NZD and buy JPY.**

$$\begin{aligned}
 \text{JPY??} &= \text{NZD1} \\
 \text{NZD1} &= \text{USD } 0.4290 \\
 \text{USD1} &= \text{JPY } 127.15 \\
 \text{NZD/JPY} &= \frac{(0.4290 * 127.15)}{(1 * 1)} \\
 &= 54.55 \text{ ASK}
 \end{aligned}$$

Again we can solve the problem algebraically, using the definition for the exchange rate:

$$\text{Ask } e_1 = \frac{\text{USD } 0.4290}{\text{NZD } 1}$$

$$\text{Ask } e_2 = \frac{\text{JPY } 127.15}{\text{USD } 1}$$

$$\text{Ask } e_3 = \frac{\text{JPY } ?}{\text{NZD } 1} = e_1 \times e_2 = \frac{\text{USD } 0.4290}{\text{NZD } 1} \times \frac{\text{JPY } 127.15}{\text{USD } 1} = 54.55$$

Thus, the two-way quote on the NZD/JPY would be 54.38-55.

There are two problems at the end of the chapter that will test your understanding of cross-rates. You may find them a little more challenging than Examples 4.1 and 4.2.

### Spot rates: various settlement periods

So far, our discussion has concentrated on buying or selling currency “spot,” which means that settlement takes place within two days of writing the contract.

In fact, there are three conventional settlement periods:

Value today	(“tod”)	Settlement today.
Value tomorrow	(“tom”)	Settlement on next business day after dealing.
Value spot		Settlement two business days from dealing date, which is the most common throughout the world. Some transactions in North America are settled in just one day.

Note that slightly different prices will apply to the rates above for several reasons which will become apparent in the next section.

## 4.5 Money Market Forward Rates

Forward transactions involve agreements where the exchange rate has been fixed, but the funds do not actually change hands until a settlement date at least three business days later.

The following information was listed in the *Wall Street Journal* on May 22, 2001. Spot currencies and forward rates for the most actively traded currencies are listed.

**Table 4.3 May 22, 2001 Section C Wall Street Journal**

Country	USD equiv	USD equiv	Currency per USD	Currency per USD
	Monday	Friday	Monday	Friday
Argentina (Peso)	1.00010001	1.00050025	0.9999	0.9995
Australia (Dollar)	0.52803886	0.52985747	1.8938	1.8873
Austria (Schilling)	0.06374868	0.06400655	15.6866	15.6234
Bahrain (Dinar)	2.65251989	2.65251989	0.377	0.377
Belgium (Franc)	0.02174523	0.0218332	45.9871	45.8018
Brazil (Real)	0.43094161	0.43468811	2.3205	2.3005
Britain (Pound)	1.4402	1.4387	0.6943	0.6951
1 Month Forward	1.4394	1.4374	0.6947	0.6957
3 Months Forward	1.4362	1.4348	0.6963	0.697
6 Months Forward	1.4324	1.431	0.6981	0.6988
Canada (Dollar)	0.65049112	0.65235828	1.5373	1.5329
1 Month Forward	0.65023734	0.65210303	1.5379	1.5335
3 Months Forward	0.64977258	0.65167807	1.539	1.5345
6 Months Forward	0.64918203	0.65108405	1.5404	1.5359
Chile (Peso)	0.00164663	0.00164663	607.3	607.3
China (Renminbi)	0.12080504	0.12081964	8.2778	8.2768
Colombia (Peso)	0.00042946	0.00042997	2328.5	2325.75
Czech Republic (Koruna)	0.02555257	0.02561082	39.135	39.046
Denmark (Krone)	0.11757098	0.11791619	8.5055	8.4806
Ecuador (US Dollar)-e	1	1	1	1
Finland (Markka)	0.14753397	0.14813059	6.7781	6.7508
France (Franc)	0.13372917	0.13426964	7.4778	7.4477
1 Month Forward	0.1336684	0.13421197	7.4812	7.4509
3 Months Forward	0.13355771	0.1340914	7.4874	7.4576
6 Months Forward	0.13345255	0.1339836	7.4933	7.4636
Germany (Mark)	0.44851094	0.45032874	2.2296	2.2206
1 Month Forward	0.44830987	0.45012604	2.2306	2.2216
3 Months Forward	0.44792833	0.44972117	2.2325	2.2236
6 Months Forward	0.4475875	0.44935742	2.2342	2.2254



**Chapter Four. FX Spot and Forwards**

Greece (Drachma)	0.00257417	0.00258475	388.475	386.885
Hong Kong (Dollar)	0.12820842	0.12820842	7.7998	7.7998
Hungary (Forint)	0.0033915	0.00341279	294.855	293.015
India (Rupee)	0.02129472	0.02130833	46.96	46.93
Indonesia (Rupiah)	0.00008722	0.00008791	11465	11375
Ireland (Punt)	1.11383382	1.11831805	0.8978	0.8942
Israel (Shekel)	0.24160425	0.24137099	4.139	4.143
Italy (Lira)	0.00045304	0.00045487	2207.3301	2198.4332
Japan (Yen)	0.00814598	0.0080952	122.76	123.53
1 Month Forward	0.00817615	0.00812338	122.307	123.1015
3 Months Forward	0.00822893	0.00817728	121.5225	122.29
6 Months Forward	0.00831514	0.00826354	120.2625	121.0135
Jordan (Dinar)	1.4068655	1.4068655	0.7108	0.7108
Kuwait (Dinar)	3.24991875	3.25203252	0.3077	0.3075
Lebanon (Pound)	0.00066039	0.00066039	1514.25	1514.25
Malaysia (Ringitt)-b	0.26315789	0.26315097	3.8	3.8001
Malta (Lira)	2.20409963	2.20799293	0.4537	0.4529
Mexico (Peso)	0.11158224	0.11076037	8.962	9.0285
Netherland (Guilder)	0.39805748	0.39966428	2.5122	2.5021
New Zealand (Dollar)	0.42700371	0.42900043	2.3419	2.331
Norway (Krone)	0.11031561	0.11046672	9.0649	9.0525
Pakistan (Rupee)	0.01616815	0.0162206	61.85	61.65
Peru (New Sol)	0.27839644	0.2777392	3.592	3.6005
Philippines (Peso)	0.01966568	0.01987084	50.85	50.325
Poland (Zloty)	0.25094103	0.25128785	3.985	3.9795
Portugal (Escudo)	0.00437545	0.00439316	228.5477	227.6265
Russia (Ruble)-a	0.03434892	0.0343678	29.113	29.097
Saudi Arabia (Riyal)	0.26663112	0.26663112	3.7505	3.7505
Singapore (Dollar)	0.55248619	0.55187638	1.81	1.812
Slovak Republic (Koruna)	0.02030189	0.02036411	49.2565	49.106
South Africa (Rand)	0.12690355	0.12618297	7.88	7.925
South Korea (Won)	0.00076805	0.00076687	1302	1304
Spain (Peseta)	0.00527208	0.00529341	189.6785	188.914
Sweden (Krona)	0.09804402	0.09806806	10.1995	10.197
Switzerland (Franc)	0.57231157	0.57388809	1.7473	1.7425
1 Month Forward	0.57273769	0.57428358	1.746	1.7413
3 Months Forward	0.57352604	0.57514235	1.7436	1.7387
6 Months Forward	0.57487784	0.57646855	1.7395	1.7347
Taiwan (Dollar)	0.02994012	0.03036745	33.4	32.93
Thailand (Baht)	0.02193223	0.02196595	45.595	45.525
Turkish (Lira)	0.00000089	0.0000009	1122500	1110000
United Arab (Dirham)	0.27225701	0.27225701	3.673	3.673
Uruguay (Peso) Financial	0.0764526	0.07655502	13.08	13.0625
Venezuela (Bolivar)	0.00139909	0.00139855	714.75	715.025
Special Drawing Rights	1.26119309	1.26326427	0.7929	0.7916
Euro	0.8772	0.8808	1.14	1.1353

Forward rates can be quoted in two ways: as simple “points” or as outright rates. In the *Wall Street Journal*, forward rates are quoted outright. When a quote is given on the basis of forward points, apply the points to the last digits of the spot rate to obtain the *all-up* or *outright* rate. Thus, suppose the following NZD/USD quotes apply with spot followed by forward points:

Foreign Exchange: Spot and Forward Points	
Spot	.4271/75
1 month	-14/-12
2 months	-29/-27
3 months	-43/-40
6 months	-84/-80
9 months	-127/-120
12 months	-167/-157

These quotes could have appeared instead as the (identical) outright forward quotes:

Foreign Exchange: Spot & Outright Forwards	
Spot	.4271 / 75
1 month	.4257 / 63
2 months	.4242 / 48
3 months	.4228 / 35
6 months	.4187 / 95
9 months	.4144 / 55
12 months	.4104 / 18

Sometimes the forward points are quoted without showing the sign. In this case, do you add or subtract the points? The rule is: if the points are *decreasing* left to right, subtract the points, and if they are *increasing*, add them. By applying this rule, the bid rate (on the left) will always be lower than the offer.

### Calculation of the forward FX price

The forward price of a currency is largely a representation of the interest differential between two currencies, not necessarily a reflection of the value of a currency at a given future date. It should reflect the break-even point of two hedging alternatives.

A company with a currency exposure several months away could either cover through the forward market or by entering a spot transaction and using the respective deposit markets, where the latter encompass borrowing as well as lending. While the actual forward price quoted by a dealer will be influenced by the dealer's existing or desired position, the range of the two-way price will be governed by the price of hedging by crossing spreads in the deposit and spot markets. We will begin with a simple mid-rate calculation in Example 4.3 and then progress to a two-way price in Example 4.4.

### Example 4.3

Sasquatch Ltd. in Canada is committed to remit USD10 million to a counter party in the US in six months. It must cover this exposure, and has decided to calculate the most effective way of doing so. It has a choice of covering with a forward rate or by entering into a spot exchange contract and using the deposit markets.

Sasquatch Ltd. finds that the current rates in the relevant markets are as follows:

Spot USD/CAD	1.5373
6 months forward margin	0.0030
6 month CAD borrowing rate	6.25% (actual /365 days basis)
6 month USD deposit rate	5.75% (actual /360 days basis)
Days in calculation	180

If the company deals in the outright forward market, the rate will be 1.5403.

The transaction will be:

The six-month forward rate is  $1.5373 + .0030 = 1.5403$ .

Sasquatch Ltd. buys USD 10,000,000 and sells CAD 15,403,000.00 at 1.5403.

Alternatively, Sasquatch Ltd. can enter into a spot transaction and use the deposit markets. Here, the company would buy USD value spot, borrow the CAD equivalent spot for six months, and lend the USD spot amount for six months.

The USD *forward* amount required would be USD 10,000,000, which must be translated into a current amount, and then into CAD.

### Solution 4.3

To compare results, calculate the implied exchange rate by:

- a) Finding the present value of the USD amount due in 180 days:

$$\frac{\text{USD}10,000,000}{(1 + .0575 \times 180/360)} = \text{USD}9,720,534.63$$

- b) Finding the total amount of the CAD loan plus interest to be repaid in six months time:

Converting USD 9,720,534.63 into CAD at the spot rate of 1.5373 gives a CAD-borrowing requirement of 14,943,377.89. In 180 days, the amount owed will be:

$$\text{CAD}14,943,377.89 \times (1 + .0625 \times 180/365) = \text{CAD}15,403,961.46$$

CAD Principal	CAD	14,943,377.89	
+ CAD interest	CAD	<u>460,583.57</u>	(365 days per year)
		CAD	15,403,961.46

Now the USD total amount owed is USD 10,000,000.

Therefore, the implied forward exchange rate is:

$$\frac{\text{CAD}15,403,961.46}{\text{USD}10,000,000} = 1.5404$$

Which rate should Sasquatch Ltd. take?

Clearly, Sasquatch Ltd. should take the quoted forward rate which is lower, because this option involves payment of a smaller quantity of Canadian dollars. It is buying US dollars, and is achieving a lower price in terms of CAD dollars.

We can summarize these calculations with the following equation:

$$\text{Forward Rate} = \text{Spot rate} \times \frac{\left(1 + i_T \times \frac{\text{no. of days}}{\text{dpy}}\right)}{\left(1 + i_C \times \frac{\text{no. of days}}{\text{dpy}}\right)},$$

where  $i_T$  is the relevant interest rate in the *terms* currency (CAD),  
and  $i_C$  is the relevant interest rate in the *commodity* currency (USD).

#### Example 4.4

A Mexican exporter will be receiving USD 10,000,000 funds in 90 days. What is the risk the exporter faces?

Suppose that the following rates prevail:

The spot rate is currently USD/MXN	8.9606 – 86
The 90 days rate in Mexico is:	8.91 – 8.98% p.a. (actual / 360 days basis)
The 90 days rate for USD deposits is:	3.92 – 3.96% p.a. (actual / 360 days basis)

Estimate the forward rate that the exporter can achieve if he wishes to hedge the above exposure. (What type of limit does this place on the outright forward rate in the market?)

#### Solution 4.4

Since the Mexican exporter will be receiving funds of USD 10,000,000 in 90 days, the risk is that the USD depreciates; that is, that the peso increases in value relative to the USD, so that the exporter receives fewer pesos in exchange for the USD earnings. The forward rate can be calculated as follows:

- a) Borrow USD @ 3.96% p.a. for 90 days. (Must borrow at the high side.)

Find the discounted amount to borrow.

$$\frac{10,000,000}{(1 + .0396 \times 90 / 360)} = \text{USD } 9,901,970.49$$

The amount to repay in 90 days will obviously be USD 10 million.

- b) Sell the USD spot at 8.9606 (i.e. the low side in terms of pesos).  
Amount of MXN to invest is:

$$\text{USD } 9,901,970.49 \times 8.9606 = \text{MXN } 88,727,596.79$$

Invest the peso at 8.91% p.a. for 90 days. (Must use low side.)

$$88,727,596.79 \times (1 + .0891 \times 90 / 360) = \text{MXN } 90,704,004.01$$

The resulting implied USD/MXN exchange rate is: 9.0704  
(90,704,004.01 / 10,000,000 = 9.0704.)

Given the above interest rates, the implied future exchange is 9.0704. If the forward rate is quoted *higher* than 9.0704, then clearly it is a good deal, better than can be achieved through the spot and deposit markets, so the exporter would likely hedge using the forward markets. If the forward rate were quoted *lower* than 9.0704, it would be better to hedge using the spot and credit markets.

One of the most difficult concepts for the newcomer to financial markets to grasp is that of a two-way price. Confusion is compounded by the fact that we look at problems from different points of view—the price-taker's and the price-maker's. It may help to remember that first of all, a price-taker always gets the worst side of the spread, and that secondly, a price-maker must price in a way that allows him or her to cross the spread in order to cover the quoted position, without incurring a loss.

- 1 a) Suppose that as a trader from Bank A, you call a trader at Bank B and ask for a quote on the Mexican peso. She quotes USD/MXN = 8.9767 - 86.

Under this scenario, what are your three options, and at what price can you transact? Note that as the caller, you will have to “cross the spread,” that is, transact at the price least favorable to you, if you do decide to transact at all.

- b) Now suppose you are an FX trader at Bank C who receives a call from an authorized trader at Bank A. You are asked for your price in Dollar/Mark (for which you are hired to deal), and you make the price USD/DEM 2.2293 - 98.

If the caller from Bank A decides to buy, at what price does the transaction go through?

2. Find the two-way GBP/DEM (British pound vis-à-vis the German mark) cross rate given the following market quotes:

Foreign Exchange Rate	BID	OFFER
GBP/USD	1.4402	1.4408
USD/DEM	2.2291	2.2296

3. Compute the bid and offer for the NZD against the CAD from the following rates of exchange. In other words, construct a NZD/CAD two-way cross rate. Here, both the NZD and the CAD are reported as commodity currencies, expressed in terms of the USD, which makes the problem a little more challenging.

Foreign Exchange Rate	BID	OFFER
NZD/USD	0.4250	0.4260
CAD/USD	0.6504	0.6510

# 4 Exercises

4. Fill in the missing legs (marked ❶ and ❷) for the no-arbitrage spot FX quote table below. Be sure to show all your calculations.

Foreign Exchange Rate	BID	OFFER
AUD/EUR	❶	0.6001
AUD/USD	.5254	0.5260
EUR/USD	❷	0.8770

- 4.5 (a) Suppose that you are a US exporter of goods to New Zealand and have a NZD receivable due in 120 days. The forward dealers are all off at a market luncheon, and their prices are not available. However, spot and interest rate dealers are at work, quoting the following rates:

Rates	Market	Two-way Price	Day Count Basis
Spot FX:	NZD/USD	0.4106 / 0.4116	
120 days interest	New Zealand	5.75 / 5.85	Actual / 365
120 days interest	United States	4.00 / 4.10	Actual / 360

What exchange rate can you assure yourself for your NZD receivable?

- (b) The forward dealers come back after a convivial lunch, and one of them quotes you 0.4068 / 0.4078 for 120 days forward. All other data remains unchanged. Can you take advantage of his computation? Explain your answer.







# Part Two

## DOING IT THE EXCEL® WAY

Part Two of this book introduces you to the Excel® spreadsheet system. Spreadsheets are an excellent way of organizing computations that repeat themselves or that require a complex layout. Consequently, spreadsheets have become a vital tool in the world of finance, and one wonders how we ever got along without them! Extensive programming skills are not required to use these tools, but you do have to remember a few basic rules and procedures. Though there are other financial spreadsheet programs, Excel – produced by Microsoft – has become the industry standard. We'll be using Excel 2000 for our discussion in this section, but most of the functions we use here can be found in other packages such as Lotus® as well.

First we shall illustrate Excel spreadsheeting principles with material on fixed interest pricing, and later we'll extend the discussion to cover complex calculations where the true advantage of using a spreadsheet will become apparent. It turns out that you can use Excel for much more than just bond pricing, and you'll find that the principles learned in this section can be readily applied in later chapters on data processing for financial econometrics and similar applications.

Those who are already familiar with Excel can probably just skim Part Two, perhaps taking a quick look at the examples, working through one or two, and trying to think up alternative ways of producing the same result. Indeed, there are often several paths to the end result with spreadsheets, and we certainly don't claim to final say in how it's done!



---

# Five

## Learning by Doing

### **An introduction to financial spreadsheets**

Spreadsheets are used extensively in the world of finance for performing calculations, analyses, and constructing charts and graphs. If you are a tertiary student or are already employed in a financial role, and find you can't deal confidently with spreadsheet programs such as Excel or Lotus 1-2-3, now is the time to brush up on your skills.

This section will walk you through a simple example from the financial application of bond pricing. Again, we'll be using the Excel package here, but similar instructions apply for Lotus.

To get into the spreadsheet program, double click on the Microsoft Excel icon on your computer. (If you can't find this icon, try double clicking on the Microsoft Office icon, and look for Excel inside the folder.)

Once Excel is open, you will see an empty spreadsheet. The spreadsheet consists of *cells*, organized into *rows* which are numbered along the left hand side starting at the top with 1 running to 66,536 at the bottom, and *columns* which are labelled along the top starting on the left with A, then B, C, and so on. Each cell has a distinct address. For example, [A1] is the address of the cell in the top left corner, with [B1] next right, and [A2] directly below. A cell can contain various types of information such as a number, words, or a formula which may refer to other cells. You can reference the contents of a cell by listing its cell address.

## 5.1 Step-by-step bond valuation example

**Task:** Find the present value on a coupon date of the following bond.

Term to maturity = 5 years  
Coupon rate = 8.0% per annum paid semi-annually  
Current market yield = 7.0% per annum (discounted semi-annually)

**Solution:**

Begin with an empty spreadsheet (or *worksheet*, as it is called in Excel).

Using the mouse pointer, locate and click on cell [A2] (first column, second row down). Type the word **Coupon** in this cell and press the **Enter** key twice:

Enter↵ Enter↵

You should now be in cell [A4]—that is, cell [A4] should be outlined in heavy black so that it stands out as the active cell.

Type the word **Principal** and press **Enter** twice again.

Enter↵ Enter↵

By now your pointer should be in cell [A6].

Type the words **Time Period** and press **Enter**.

Enter↵

Oops! By now some of you may have made a small mistake. For instance, if you typed “Time” instead of “Time Period” by accident and hit the enter key, you can correct the problem without much difficulty. Use the up arrow key to return the pointer to [A6] (or just click on the cell [A6] again). Then press the F2 key to get into edit mode, and fix your typo. Then simply press the enter key again.

Enter↵

Excel recognizes the words you have typed as text (or labels) and treats them accordingly. For calculation purposes, Excel will assign a zero value to any cell that contains text.

“Time Period” will not fit neatly into the cell, so run your mouse pointer up to the top of the column to the label row of letters just above row 1. Center the mouse on the dividing line between column A and B. Click and (holding down the mouse key) drag the line to the right, causing row A to become wide enough to fit the words into cell [A6]. (Alternatively, simply double click on the dividing line between A and B at the top of the column, and the column width will change automatically to accommodate the label “Time Period.”)

Next, click on cell [B2] and type in **8%**; then press **Enter**.



It would be helpful to show two decimal places, so click again on cell [B2]. Run the mouse up to the formatting tool bar (the second tool bar above your worksheet), and click twice on the *increase decimal* format button that looks like this:



Cell [B2] should now read as 8.00%.

Click on cell [B4] and type in **1000000** next to the word **Principal** that sits in cell [A4]. Press **Enter**. As it stands, the number you have typed into [B4] is a bit difficult to decipher at a glance. (Is it 10 million or 1 million?) Return the cell pointer to [B4] by using the arrow keys or by clicking on [B4] with the mouse, then use the mouse to click on the comma button in the format bar:



The number in [B4] will become more readable as 1,000,000.00

Note that you can type in the commas as you enter the number, in which case Excel will accept them as an indication of how you wish all numbers to be presented. Similarly, because you specified above that you wanted 8.00% written with two decimal places, Excel will “remember” this request and assume that you want two decimal places for other values on the same worksheet.

With commas and two decimals, the number in B4 becomes too large to fit into the current cell width. Excel 2000 will automatically make the column wider to accommodate your number. Some older versions of Excel will fill the cell with a series of hash marks, #####, indicating overflow. For versions prior to '97, you

will have to increase the width of column B as you did for column A. Click on the dividing line between labels B and C at the top of column B and drag it to the right until there is enough space to show 1,000,000.00 in [B4]. (Alternatively, double click on the dividing line between labels B and C and the column width will adjust automatically.)

Now we are going to number the ten semi-annual time periods during which the bondholder will receive coupon payments. Click on cell [A9] to make it active, or activate it by moving the cell pointer there using the arrow keys.

Type the number one (1) into cell [A9], and press **Enter**.

Enter↵

Type the number two (2) into [A10], three (3) into [A11], four (4) into [A12], and so on, through ten (10) in cell [A18], pressing **Enter** after each entry.

Enter↵

Column A should now contain the numbers 1 through 10 running from cell [A9] to [A18].

Click on cell [B6] and type the words **Cash flows**. Then press **Enter** three times.

At this point, you should have your cell pointer in [B9]. In this cell, you are going to place a formula to calculate the bond's first cash flow. This is how it's done:

- (1) First, to check that you are in cell [B9], look at the formula bar which lies immediately above the column labels (A, B, C, etc.) on your worksheet. You will see the current cell address in the left-most section of the formula bar. Follow the rest of the instructions below, watching what appears in the formula bar.
- (2) Type the arithmetic operator, = (the *equals* sign), into [B9], and then click on cell [B4]. (Note that you do not press the enter key yet.)
- (3) Next, type in the arithmetic operator, \* (shift-8), which stands for multiply, and click on cell [B2] (still avoiding the enter key).
- (4) Type the arithmetic operator, / , which stands for divide, followed by 2, and click the enter button in the formula bar which appears as a green tick. The



cell pointer should remain in [B9]. In the formula bar, to the right of the cell name (B9), there should now be the formula **=B4\*B2/2**.

Two points are worth mentioning at this juncture:

- (1) To finish the formula, you could have pressed the standard enter key or one of the arrow keys, but then the cell pointer would have moved one cell down or one cell in the direction of the arrow.
- (2) You could also have simply typed the formula =B4\*B2/2 into cell [B9] without clicking on cells [B4] and [B2], but the pointing method described above is both safer and faster, once you get used to it.

Excel recognizes that you have entered a formula in cell [B9], and responds by performing a calculation and showing the answer—40,000—on the worksheet. When the cell pointer is on [B9], the formula will show in the formula bar, but on the worksheet itself you will only see the answer, 40,000 in this case, showing in the cell.

### **Relative and absolute addressing**

The formula in [B9] is fine as it is for that cell, but now we want to copy the formula and paste it into the five cells immediately below [B9]. We want Excel to perform the calculation of the interest cash flow as before, and report the answer as 40,000.00 in each of the cells.

Unfortunately there is a snag. Whenever you enter a formula, Excel assumes it is in *relative address format*. For many copy situations, this format would be fine, but not in this case. If we copy the formula as it is written into cell [B10], the new formula in [B10] will read: **=B5\*B3/2**. Excel will adjust the formula by the same amount and direction that you move your cell. The result on your spreadsheet will read zero (0), which is clearly not the cash flow we want.

Hence, it is necessary to change the formula in [B9] into *absolute address format*, which will anchor the cell references in the formula to the cells [B4] and [B2], so that when we copy the formula those references will stay the same.

We can edit the formula in [B9], changing the cell references to absolute addresses in the following way.

Double click on cell [B9]. Using the arrow keys, move the flashing indicator within the formula until it is between “B” and “4” (in other words, zero in on the [B4] cell reference).

Then press the **F4** key (top row of the keyboard).

**F4**

The cell reference will change to read **\$B\$4**, indicating it is now absolute. Move the indicator to the [B2] reference in the formula and press F4 again. The [B2] should change to **\$B\$2** and all cell references in the formula will now be *absolute cell references*. They will not change when the cell is copied.

Click the green tick ✓ in the formula tool bar.

### Copying a formula

We can now safely copy the contents of [B9] into cells [B10] through [B18] as follows.

- (1) Make sure the cell pointer is in [B9].
- (2) Click the copy button in the standard toolbar (the icon depicting two sheets of paper just to the right of the printer button). To be sure it is the copy button, position the mouse over the button before clicking and the name of the tool will appear.
- (3) Highlight cells [B10] to [B18] by clicking on cell [B10] and dragging the wide cross mouse indicator through to cell [B18].
- (4) Click the paste button (to the right of the copy button, with a clipboard and paper).

Alternatively, if you are a confident rodent controller, you could do the above copy by clicking on [B9], then holding the left mouse button down, and with the small cross (+) indicator showing at the lower right corner of the cell, dragging the mouse to [B10] through [B18].

If you feel like giving the second copy technique a try, start by pressing the delete button (marked *Del*) to remove the highlighted data you have just entered using the paste procedure. Then:

- (1) Click on cell [B9] and release the mouse button. Now position the mouse on the lower right corner of the cell (the mouse pointer will appear as a + sign to the right).
- (2) Hold down the left mouse button and drag the outline of the cell you have selected into [B10] through [B18] and release the mouse.

Now you have copied the formula from cell [B9] into cells [B10] through [B18]. Click on any of these cells and view the formula bar. It should show  $=\$B\$4*\$B\$2/2$ . On your spreadsheet, it will show as the value 40,000.00. This is the result we want for all of these cells except the last one, [B18]. Here we must use the editor again.

Double click on [B18] to edit, or get into edit mode by clicking on [B18] once and pressing the **F2** key.

Extend the formula by adding **+B4** at the end of the formula to include the repayment of the principal. Use the arrow keys to move within the formula.

Press the **F4** key to make **+B4** an absolute reference.

**F4**

Click the green tick  $\checkmark$ .

All the cash flows should now be correctly represented in cells [B9] to [B18].

### Calculating the discount factors

The next step in our bond spreadsheet involves calculating the discount factors to apply to the series of cash flows to obtain their values in today's terms.

- (1) Click on cell [C2] and type in the word **Yield**. Press **Enter**.

- (2) Click on cell [D2] and type in the yield to maturity as 7%; press **Enter**.
- (3) Click on [C6], type the word **Discount**, and press **Enter**. Type the word **Factor** into cell [C7] immediately below.
- (4) Click on cell [C9]. Type the following keystrokes:  $=1/(1+$
- (5) Click on [D2] (which is where the discount rate or yield to maturity is stored), and press the **F4** key.
- (6) Type:  $/2)^$  (the last character shift-6).
- (7) Click on [A9], but do NOT press F4 this time.
- (8) Click on the *enter formula* button, which appears as a  $\sqrt{\quad}$  in the formula tool bar.

At this point, the formula in [C9] should read  $=1/(1+\$D\$2/2)^A9$ . The “^” symbol means *raise to the power of*, and the “A9” refers to the number of time periods for which you will be discounting.

As the first cash flow arrives after one time period, you must discount for just one period. However, as the second cash flow arrives at the end of two time periods, you will want it to be discounted back two time periods. Leaving the cell reference to [A9] as a *relative reference* will ensure the appropriate discounting when the formula is copied.

Click on cell [C9] and copy as before into cells [C10] through [C18]. You have now completed the calculation of the ten discount factors.

### Applying the discount factors

We will now use the discount factors we have calculated in cells [C9] to [C18] to value the cash flows in [B9] through [B18], and place the results in column D.

- (1) Click on [D6]. Type the phrase **PV of** and press **Enter**. Then enter the words **Cash flows** into cell [D7] below.

- (2) Click on [D9]. Type in = (the *equals* sign). Don't touch the Enter key yet.
- (3) Click on [B9]. Then type \* and click on [C9]. (Don't touch the F4 key, as we want all these address references to remain relative.)
- (4) Click on the *enter formula* button ( $\surd$  in the formula tool bar).

Copy the formula in cell [D9] into cells [D10] through [D18]. Then click on cell [D19]. In this cell, we are going to put the total present value of the bond. Run the mouse up to the standard tool bar and click on the *sum* sign, which is depicted by the sigma symbol:



Check to see that it is summing cells [D9] to [D18]; if it doing so correctly, press **Enter**. (If it is not summing correctly, you may need to edit the formula.)

Your spreadsheet is now almost complete. However, it could do with some formatting help to tidy it up and make it easier to read.

Highlight the cells [B9] to [B18]. Hold down the *control* key and highlight the range of cells [D9] to [D18]. The control key acts as an indicator to select the second range as well as the first. Both cell range selections should appear highlighted. Click on the comma button in the format tool bar.



All the highlighted cells will now be formatted with commas.

Finally, click on cell [D19] and then run the mouse up to the formatting tool bar again to click on the dollar sign to the left of the comma.



This action puts a dollar sign in front of the number and helps to emphasize that it is a total. (Recall that you may have to adjust the column width for older versions of Excel.)

## Printing and saving

Now we will print your masterpiece, but before we do that, we need to save it.

To save for the first time, run your mouse up to the menu bar which holds the various Excel pull-down menus such as *File*, *Edit*, *View*, etc. Click on the *File* menu and select the *Save As...* command (by clicking once again). A dialog box will open (the three dots, or *ellipsis*, always indicates that a dialog box accompanies that selection).


Type a suitably descriptive name in the *File name* box, so that you'll remember what this file contains when you come back to it in the future.

You may want to keep your files organized by saving your spreadsheet on a drive or in a folder other than the default; if so, click the drop-down button for the *Save in* list box, click on the disk drive you want to use, and then click the appropriate folder.

When you are building your worksheet, it is a good idea to save your document often. That way, in the event of a power failure or an inadvertent keystroke, you won't lose too much of your work. Just click the *File Save* button in the standard tool bar. (It appears as a floppy disk button about third from the far left of the worksheet.)

## Prettying up for printing

For printing purposes, it is often helpful to place a border around a table in a worksheet. Highlight the entire table from cell [A2] to cell [D19]. Run the mouse up to the format tool bar and click on the *Borders* drop-down list button. (The icon depicts arrow pointing down next to a box.)



Select a border from the options displayed. Alternatively, you can select *Format* from the menu toolbar, then *Cells...*, *Border*, and make your choice from the dialog box.

Leaving the highlighting in place, run the mouse up to the standard tool bar and click on the print button. (The icon looks like a printer and is next to the *Save* button).

Your first bond spreadsheet is now complete. Printed, it will look something like Figure 5.1.

## 5.2 Do-it-yourself fixed interest workshop

### Bonds: calculating their prices and yields

Bonds are longer-term securities with original terms to maturity of greater than one year. For example, auction data from the June 29, 2001 issue of the *Wall Street Journal* showed that U.S. Treasury bonds range from maturities of June 30, 2003 to February 15, 2031.

Coupon	8.00%	Yield	7.00%
Principal	1,000,000.00		
Time Period	Cash flows	Discount Factor	PV of Cash flows
1	40,000.00	0.966184	38,647.34
2	40,000.00	0.933511	37,340.43
3	40,000.00	0.901943	36,077.71
4	40,000.00	0.871442	34,857.69
5	40,000.00	0.841973	33,678.93
6	40,000.00	0.813501	32,540.03
7	40,000.00	0.785991	31,439.64
8	40,000.00	0.759412	30,376.46
9	40,000.00	0.733731	29,349.24
10	1,040,000.00	0.708919	737,275.57
			\$1,041,583.03

**Figure 5.1 Pricing a five-year bond**

Recall that the basic features of a bond include:

1. Its denomination or **face value**; for example, \$1,000,000.
2. Its **maturity date**; for example, June 30, 2003.

3. Its **coupon rate**; for example, 3.875% p.a. (1.9375% semi-annually).
4. The **frequency** of coupon payments each year and the specific dates of each payment.
5. The **identity** of the issuer.

For U.S. Treasuries, interest is paid semi-annually from the date of issue. The date of issue is usually the 15<sup>th</sup> of the auction month (though for some maturities, it is the end of the month). Auctions are generally monthly for shorter maturities, and quarterly for longer maturities.

### Bond pricing

A bond's value is determined by the present value of its future cash flows. These cash flows are:

- the stream of fixed coupon payments, and
- the repayment of the bond's face value at maturity.

Recall the equation for the price of a coupon bond (P), Equation (9) from Section 2.6:

$$P = \frac{C}{1+i} + \frac{C}{(1+i)^2} + \dots + \frac{(C+FV)}{(1+i)^n} \quad (1)$$

where

C	=	regular coupon payment
i	=	interest rate per compound period
FV	=	face value of bond
n	=	number of interest periods (generally use half years).

### Problem 5.1

Find the price of a three-year bond of face value \$1,000,000, paying 8.0% coupons on a semi-annual basis, given the following interest rate scenarios.

Use Equation (1) above and a spreadsheet table for pricing the bond at a market yield of 7.0% p.a. (a format for the table is suggested below).

Start by calculating the cash flows from the bond in the second column, remembering to include the repayment of principal at the end of the sixth period. Then compute



the discount factors in column three. Multiply each cash flow in column two by its discount factor in column three, placing the result in column four.

The total price of the bond is given by summing column four to obtain the total of the PV's of the cash flows.

**Table 5.1 Pricing a three-year 8.0% bond at yield of 7.0%**

<b>Coupon:</b>	<b>8.00% p.a.</b>	<b>Yield:</b>	<b>7.00% p.a.</b>
<b>Principal:</b>	<b>\$1,000,000</b>		
<b>Time Period</b>	<b>Cash flow</b>	<b>Discount factor at yield of 7% p.a.</b>	<b>PV of cash flows at yield of 7.0% pa</b>
1			
2			
3			
4			
5			
6			
Total			

(Note that if you seldom use spreadsheet programs, you will probably find it helpful to reference Section 5.1, which provides a step-by-step solution to a similar problem.)

When setting up the spreadsheet, remember:

- The bond pays a *semi-annual* coupon.
- The yield is quoted on a nominal annual basis, and it represents semi-annual compounding. What discount rate will you use for a nominal annual yield of 7.0%?
- For calculation purposes, yields can be expressed as decimals or in the percent format (e.g., 9.0% or 0.0900—Excel will accept either.)

Then price the same bond at a yield of 9.0% p.a. You can copy the table to create another one for pricing at 9.0%, or you could just add two more columns to the one you've already created.

**Table 5.2 Pricing a three-year 8.0% bond at yield of 9.0% p.a.**

Coupon	8.00% p.a.	Yield	9.00% p.a.
Principal	1,000,000		
Time Period	Cash flow	Discount factor at yield of 9.0% p.a.	PV of cash flow at yield of 9.0%
1			
2			
3			
4			
5			
6			
Total			

**Question:**

Having found the price for the bond at a 7.00% yield and at 9.00%, what relationship have you demonstrated between the price of the bond and the market yield?

A bond's coupon payments form an annuity; thus, the value of any bond is given as the present value of this annuity added to the discounted amount of the face value.

Using the *annuity* formula (see Section 2.6), we can short-cut the lengthy calculations as follows:

$$P = C \times PVIFA_{i,n} + \frac{FV}{(1+i)^n}$$

where the PV of an annuity for  $n$  periods at rate  $i$  is:

$$PVIFA_{i,n} = \frac{\left(1 - \frac{1}{(1+i)^n}\right)}{i}$$

**Problem 5.2**

Find the price of a ten-year bond of face value \$1,000,000 paying 8.0% p.a coupons on a semi-annual basis, given the following interest rate scenarios. This time, build your spreadsheet using the annuity formula. You can probably build the three interest rate scenarios into the same spreadsheet, since the annuity formula will shortcut the number of lines in the calculation considerably.

Market yields:

- (a) 7.0%
- (b) 8.0%
- (c) 9.0%

A possible layout for the table is seen in Table 5.3.

**Questions:**

1. On the ten-year bond in Problem 5.2, what did you notice about the price of the bond when the yield was 8.0%?
2. Try adjusting your spreadsheet from Problem 5.1 by substituting an 8% yield (or better yet, adding extra columns to incorporate this sensitivity). What is the result?

**Table 5.3 Pricing a ten-year 8.0% bond at yields of 7.0%, 8.0% and 9.0% p.a.**

<b>Coupon:</b>	<b>8.0%</b>	<b>Principal:</b>	<b>\$1,000,000</b>	
<b>Yield</b>	<b>Coupon flow</b>	<b>PV of coupons</b>	<b>PV of principal</b>	<b>Total PV</b>
7.0%				
8.0%				
9.0%				

3. Can you begin to form a basic proposition from this result and the one above for the ten-year bond?

### Bonds at a premium, discount or par

A bond whose price (P) exceeds its face value (FV) is said to be priced at a *premium*. A bond whose price is less than its face value is priced at a *discount*. If a bond's price is equal to its face value, then the bond is priced at *par*.

### Problem 5.3

Make a summary table like the one below using the results of the previous two problems and marking the last column P (premium), D (discount), or Par (face value = price).

**Table 5.4 Summary of results for problems 5.1 and 5.2**

Bond	Coupon	Yield	Face Value	Price	P, D, or Par
3-year bond	8.0%	7.0%			
	8.0%	8.0%			
	8.0%	9.0%			
10-year bond	8.0%	7.0%			
	8.0%	8.0%			
	8.0%	9.0%			

**Table 5.5 Coupon, yield and bond price**

	Price vs. face value	Relationship of C to $i$
Discount	$P < FV$	
Par	$P = FV$	
Premium	$P > FV$	

### Bond Pricing: Relationship between coupon rate (C), yield (i), and price (P)

Now you can use the results above in Table 5.4 to summarize the relationship between the bond's coupon, the market yield, and the price of the bond. You could set out your findings in a table like Table 5.5, filling the third column with the coupon and yield relationship (that is, C is either =, >, or < i).

Note that the relationships you have revealed only hold when the bond is priced *on a coupon date*. When there is less than six months to run until the next coupon is due, a bond's price may be greater than its face value, even though the coupon rate is the same or even less than the yield.

### Finding the IRR given the price of a bond

In most of the world's interest rate markets, bonds are quoted on a *price*, rather than yield basis. If we are given a bond's price, how can we calculate its IRR or yield to maturity?

#### Problem 5.4

Find the yield to maturity for a two-year bond paying a semi-annual coupon of 8.5% p.a. when the bond is quoted at a price of 106.35 per \$100 face value.

Using a spreadsheet program such as Excel, we can set up a table similar to the one we used in Problem 5.1 and employ one of Excel's built-in financial functions—

**Table 5.6 Finding the IRR**

Coupon:	8.50% p.a.	<b>Yield:</b>	<b>? % p.a.</b>
Principal:	\$1,000,000	Target Price:	106.35
<b>Time Period</b>	<b>Cash flow</b>	<b>Discount factor</b>	<b>PV of cash flows</b>
1			
2			
3			
4			
<b>Total</b>			Total PV of Cash flows
		<b>PV per \$100 face value:</b>	

IRR. It will perhaps be more instructive, however, to build our table as before with just two small changes: a cell near the top to contain a target price, and a cell at the lower right to calculate current price per \$100 of principal.

We can adjust the yield input (see the last column on our table) until the total present value of the bond's cash flows per \$100 of principal is equal to the "target" price. The advantage of building the spreadsheet this way is that we can readily see the relationship between price and yield. When you start inputting a guesstimate for yield, will your first attempt be above or below 8.5%?

Note that if you would like to try Problem 5.4 using Excel's IRR function, you will have to add another cash flow for time period 0, the settlement amount you pay out for the bond (remember that the price is the amount you are paying to receive a set of future cash flows. So enter the price as a negative and all other cash flows as positives). You won't need to calculate the discount factors or PV's of cash flows as Excel will do that (invisibly) for you. Using the IRR function will return you the yield to maturity on an *effective semi-annual basis*. You will have to multiply the result by two to obtain the *nominal annual yield*, that is, to get back to the standard way that bond yields are stated.

### **Time out for some market jargon...and a useful concept: the PVBP**

When you performed the calculations for Problem 5.1, the yield was moved by one percent at a time. For example, you calculated the price of the bond at 7.0% and then at 8.0%. This one percent change comprises 100 basis points. If rates move, say, from 7.0% to 7.10%, we describe this as a yield increase of ten basis points. A rate increase from 7.10% to 7.11% is a move of just one basis point.

For purposes of calculation, we frequently represent yields in decimal format, so that 7.10% appears as 0.0710 and 7.11% is written as 0.0711. In other words, 0.0001 represents one basis point.

*The price value of a basis point (PVBP)* is the measure of the change in the price of a financial instrument such as a commercial bill or a bond (or a futures contract, etc.) when its market rate of interest moves by one basis point.

**Problem 5.5**

- (a) You have already priced the three-year bond in problem 5.1 at 7.0% p.a. Try increasing the yield by one basis point and pricing again. That is, find the price at 7.01% p.a.
- (b) Find the difference between the values at these two yields. This difference is the PVBP of the instrument at a yield of 7.0% p.a.
- (c) Try the same exercise with the ten-year bond. What do you notice about the PVBP this time?
- (d) Compare the PVBP's of the three-year and ten-year bonds. What relationship can you hypothesize?

**Bond Pricing – less than one discounting period from the next coupon (See Section 2.6)**

In the real world, you typically do not buy a bond at one second after a coupon anniversary date, so it is usually necessary to use fractional discounting when pricing bonds. Section 2.6 discusses this issue further. Equation (11) from Chapter 2, employing fractional discounting, is reproduced below as (2).

$$PV = \left( \left( \sum_{k=0}^{k=n} \frac{C}{(1+i)^{\frac{a}{b}+k}} \right) + \frac{FV}{(1+i)^{\frac{a}{b}+n}} \right) \quad (2)$$

A full description of the terms of the formula is provided in Chapter 2.

Remember that the fraction  $\frac{a}{b}$  allows for part period discounting.

For Problem 5.6 below, you will need to use the closed form for (2) which is reproduced as Equation (3) below:

$$PV = \left[ (1+i)^{-\frac{a}{b}} \right] \times \left[ C \times (1 + PVIFA_{i,n}) + FV \times PVIF_{i,n} \right] \quad (3)$$

Further complicating the issue of bond pricing is the fact that Equations (2) and (3) provide the invoice price, but not the *quoted price*. However, these equations will suffice for the next problem, which asks for invoice price.

**Problem 5.6**

If time permits, try to find the settlement price of the November 15, 2006, T-Bond stock which carries an 8.0% p.a. coupon when the current market yield is 4.04% p.a. Use September 24, 2001 as the settlement date.

- (a) What are the semi-annual coupons per \$100 face value?
- (b) How many full coupon periods are left to run?
- (c) How many days are there from settlement date to the next interest date?
- (d) Calculate the number of days from the last interest date, May 15, 2001, to the next interest date, November 15, 2001.
- (e) What is the yield applied for discounting?

You will need to answer these questions before you can price the bond.

For a challenge, try to put this question onto a spreadsheet. The hardest parts to program in a general format are (b), (c), and (d).

If you want to get very professional, note that bonds become ex-interest ten working days before a coupon date. You could include a calculation for whether the bond is ex- or cum-interest for the settlement date specified.







# Part Three

## STATISTICAL ANALYSIS AND PROBABILITY PROCESSES

This section of the book is about statistics and probability. Roughly speaking, statistics is concerned with the analysis of data, and probability with what we should expect the data to reveal, based on an underlying model of how the world works. Financial econometrics represents a fusion between the two. Also, the modern theory of derivatives pricing utilizes models of asset prices or returns that follow dynamic paths through time that incorporate random behavior. To understand these models, we have to know a bit about stochastic processes, which again incorporate some basic ideas about probability.

This part of our book attempts to put all these concepts on an operational footing. Needless to say, a lot of reading can be done on the subject, from basic texts on economic statistics to journal articles on the latest technological gimmicks. Even books on elementary statistics for business and economics cover statistical theory in more depth than we will here. However, some statistics textbooks lack some of the important interpretive dimensions for financial management, and as a result, there is still much misunderstanding among practitioners and academics about rather important subtleties. Some of these subtleties would take extensive explanation—not to mention technical background—but the gist can be given in more or less everyday language.

Before we proceed, however, another word in preparation. From now on, you are supposed to be fairly conversant in Excel, and this expertise will now be integrated with the text, rather than being set apart. However, we will still provide step-by-step instructions for the implementation of various data manipulations. In addition to illustrating the quantitative techniques, the information will increase your fluency in Excel as well.

To produce histograms like those shown in Chapter 6.3 and perform regression analysis as we did in Examples 7.3, 7.4, 7.5 and 8.1, it will be necessary to have the Excel Analysis Toolpak (a standard Excel add-in) up and running on your computer. (From the *Tools* menu, select *Add-Ins...* and click on the check-box for *Analysis Toolpak*.)



---

## Statistics Without Becoming One

This chapter is about statistics, though it does introduce the reader to some basic ideas on probability as well. Primary emphasis, however, is on data exploration and the principles that govern what we can reasonably infer from the data, an area known as *statistical inference*. We shall also take this opportunity to introduce you to the art of downloading data in a format that you can use for your statistical number crunching. You may find this useful even if you have no intention of doing statistics, but simply need the data for other purposes.

## 6.1 Notation and definitions

Density function	The frequency histogram, or in theory, a function where the area beneath the curve and between two points on the horizontal axis represents the probability of getting values between those two points.
Distribution	A general term applied to either the distribution function or the density function of a random variable or set of random variables.
Distribution function	The cumulated frequency histogram or density, a function whose height at any point is the probability of getting values of the random variable less than or equal to that point. Sometimes we talk about an empirical vs. a theoretical distribution function. The difference is that the former is derived from a sample taken from the population while the latter is derived from the population.
Generation process	Underlying real world model that generates the observed data, especially in a dynamic context.
Kurtosis	A measure of whether the shape of a given probability or frequency density is flatter than a comparable normal density shape. The distribution is <i>leptokurtic</i> if its peak is sharper and its tails are fatter than the normal density. The opposite is called <i>platykurtic</i> .
Mean	The arithmetic average, or in theoretical work, the sum of the variables' values weighted by the probabilities of getting those values (in this context called the expected value or expectation).
Moment	The average or expected value of the deviations from the mean raised to any designated number, usually a positive integer. The mean is the first order moment and the variance is the second order (power two) moment.
Population	The assumed parent set of observations from which the sample is drawn.

Random sample	A sample drawn from the parent population according to some process which assigns each observation an equal probability of being drawn (technically, a simple random sample).
Sample	A subset of observations taken or drawn from a parent population which itself may be too large, expensive or otherwise difficult to study, usually by means of some sort of randomization device to ensure that it is a representative sample.
Skewness	The extent to which a frequency histogram or probability density is asymmetric. The density is positively skewed if it has a longer tail to the right than to the left.
Standard deviation	A statistic used as a measure of the dispersion or variation in a distribution. Equal to the square root of the average squared deviations from the mean.
Statistic	A number computed according to some rule from the set of observations on one or more random variables.
Variance	The square of the standard deviation, or the average sum of squared deviations from the mean.

## 6.2 Introduction to data

Data is where the story begins for this chapter. It comes in every shape and form, and can generally be divided into three categories: good data, not-so-good data, and truly bad data. Before you spend any time working with a set of data, it is imperative that you first consider whether it is appropriate for the intended purpose. If you are in the position to generate your own data, it pays to spend some time (and perhaps money) determining exactly what your purpose is and what kinds of claims and inferences you need to be able to draw from your data. For example, if you have survey data, you might ask whether it is free from selection biases. Questionnaires are notorious for containing bias; even where the response rate is good, there are issues of neutrality with respect to the dimensions being studied.

For our purposes in this chapter, we will sidestep these concerns and use data provided specifically for this text on the Authors Academic Press (AAP) website (see [www.AuthorsAP.com](http://www.AuthorsAP.com)). On the AAP website, you will find a database of

monthly returns in two files called “USdata.txt” and “USdata.xls.” Don’t worry yet about trying to open them. The files contain monthly returns on the following:

1. A set of 14 large companies of which 13 are quoted on the New York Stock Exchange and one, namely York Research, is on the NASDAQ. (Not necessarily a representative set.)
2. The S&P500 Gross index.
3. The US 30-day CD rate (reported on a monthly effective yield basis to be more directly comparable with the other returns given).
4. Gross Return on US ten-year Treasury Bonds.

The data is nicely lined up over the period March 1996 – January 2001, to give 59 observations, for each of which the returns on stocks, market index, CD rate, and bonds refer to the same number of business days. This type of data is referred to as *non-overlapped*, because the figure for each month is the return over that month, with no contribution from the previous month. You could, for instance, have constructed a much longer sequence over the same span by adopting a 30-day measurement frame that moves along one day at a time. This would result in *overlapped* data, as successive observations would share, say, 28 days each.

The returns data is constructed in terms of what are called *accumulation* indexes, of which they are percentage changes. This is a way to allow for the effects of the elements of reward from holding a stock—i.e. things like dividends, rights issues, bonuses, stock splits and so forth. The basic idea is that you start with one unit of the stock and plough back all the reward elements into purchases of further units, at market prices: the accumulation index is the value of the fund that you end up with at the stated point in time. The S&P500 Gross index is this kind of accumulation index.

### 6.3 Downloading data

Downloading data should in theory be an easy task, but in practice you will sometimes run into difficulties because your data source, whatever it is, may not be in Excel format. If this is the case, you can frequently download it in an alternative format called a text file, which corresponds as closely as you can get to ordinary typing format. What we shall do in this section is to give you some directions for downloading from a text file. Downloading from an Excel file is straightforward, so if you think that you will never have to use text file format just skip Section 6.3. You may find our treatment of the text file downloading useful for other purposes:



For example, some of the older statistics packages do not allow you to feed in data as an Excel file, but do accept text files (often written with the suffix \*.txt). Downloading or feeding in via text is a bit of a hassle because it will not easily allow you to feed in variable names as well as the actual data (easy with Excel). So you have to keep on your toes a bit.

As usual we shall follow a learning by doing path. The first thing to do is examine the US data files on the AAP website. You will see that there are two files: “USdata.txt” and “USdata.xls.” We will pretend that the second file does not exist and create it from the first, which is a text file. You may notice another file called “Names.xls” that holds the names and descriptions of the 14 companies, in order of their appearance in the data.

The data comes in the form of a matrix. Each line represents an observation (one time point) on all the variables, the first being the CD rate, the second the S&P500 (or *market rate*), the third, the 14 companies, and finally, the return on the Bond index. Thus, the data is presented in a matrix of 59 rows and 17 columns. Each column represents the entire set of observations on one of the variables. Each row is a snapshot taken at some point in time of the returns on all the variables.

If you download the “USdata.txt” file, it will appear on your directory as an icon marked “USdata” which looks like a spiral-bound notebook. By double clicking on the icon, the contents of the data file can be displayed. The file has been created as a text file, which means that the observations have been entered one by one, with a tab between each figure to “delimit” it—that is, to set it apart from the ones around it—and a hard return at the end of each line. Text files are a good way to capture data if you are not sure what software package will ultimately be used to process it, because most packages will recognize a text file. We are going to convert this data to an Excel file here, nicely labelled and formatted for readability (just like the one on the website, but don’t cheat just yet!). Here’s how to proceed.

Left click on the Excel icon to start up the program. Excel will produce a fresh workbook with the title bar at the very top showing “Book 1.” Check that the cell pointer is in the top left corner of the worksheet on the cell [A1], then run the mouse pointer up to the menu bar and click on *File*, then *Open*. This will generate a dialog box which you can use to import the “USdata” text file. Beginning at the top left of the dialog box, click the drop down button to the right of the drop down list marked *Look in:* and click on the directory where you have saved the text file. Then at the bottom left of the dialog box, click the drop down button by the list marked *Files of type:*, and select *All files*. Next, select the “USdata” notepad file from the large list box above, and finally, to complete your selection and close the dialog box, click the *OPEN* button near the bottom on the right side.

If these steps have been executed correctly, pressing the open button will have started up a new routine called the *Text Import Wizard*, with a series of three step-by-step dialog boxes.

For Step 1, click on the radio button, *Delimited*. Then check that the next box indicates *Start import at row 1*. If this box does not show “1”, then using the mouse, click on one of the spinner buttons to the right, using the down arrow to decrease the number to 1, or the up arrow to increase the number to 1. Select *Windows (ANSI)* from the list box to the right and press the *NEXT* button at the bottom.

For Step 2, select  $\surd$  *Tab* as the delimiter, and click on the *NEXT* button near the bottom of the dialog box.

For Step 3, select the *General Column Data Format*, and then select *Finish*. This last step will cause the data to be imported to your new workbook. At the same time, a new name, *USdata*, will appear as the workbook title. At this point, select *File* and then *Save As...*, from the menu bar, and give your newly created workbook a new name of your own choice—something *other* than “USdata” to avoid overwriting the existing Excel file. Click on the drop-down button for the list *Save as type:*, and select *Microsoft Excel Workbook*. Finally, click the *Save* button at the bottom left.

## 6.4 Data exploration

OK, now we have all sorts of data. Eyeballing it won’t tell you much, and certainly won’t provide a basis for any claims, so let’s try a bit of exploratory data analysis. To make the data handling more meaningful, it will help to begin by putting labels at the top of your data columns. (If you skipped the previous downloading section and have opted to use the data as provided in Excel format, then you can skip the following two sentences as well. The Excel data file from the AAP website is already labelled.) With the cell pointer at [A1], use the mouse pointer to click in the menu bar on *Insert*, *Rows*. Then type abbreviated names in the new first row starting with **CD rate** in column A, **Market** in column B, **IBM** in column C, and so forth (taking the names in order of appearance from the file “Names.xls.”) The bond returns appear at the end of the data in column Q.

Pick a stock at random—say, IBM—which occupies Column C of your data. Do a *histogram*, which is a way of summarizing the empirical distribution of returns for this stock. We divide the horizontal axis into a number of intervals, and look at the

number of times that a return falls inside each chosen interval. Dividing this by the total number of observations (59) gives the relative number of times, expressed as a frequency. This is plotted on the vertical axis; the resulting construction appears as a vertical bar chart. Of course, the whole thing will be done for us by the Excel program, and here is how it is executed.

With your Excel workbook containing the returns data open in front of you, run the mouse pointer up to the menu toolbar, select *Tools*, then *Data Analyses....* Again, note that the three dots indicate that there is more to follow—in this case, a drop down list box of analysis tools. Select *Histogram*, and click on *OK*. This selection will cause the generation of a histogram dialog box. (This seems long and complicated at first, but is actually very easy when you get the hang of it.)


Respond to the histogram dialog box as follows:

- |   |  |
|---|--|
| <i><u>I</u>ntput range:</i>   | Using the mouse pointer, click and drag to highlight the third column of your data from row 1 to row 60.   |
| <i><u>B</u>in range</i>   | Skip this box, as the histogram routine will produce its own equally-spaced intervals with far less difficulty than you can.   |
| <input checked="" type="checkbox"/> <i><u>L</u>abels</i>                | Click the check box and a tick will appear indicating that option is selected.   |
| <input type="radio"/> <i><u>O</u>utput range</i>                        | Click the round “radio button” and type the cell name of the location on the spreadsheet where you wish to place the histogram. Choose a cell well clear of the data, such as [\$AA\$2]. |
| <input checked="" type="checkbox"/> <i><u>C</u>umulative Percentage</i> | Click on <i>Cumulative Percentage</i> to select.   |
| <input checked="" type="checkbox"/> <i><u>C</u>hart output</i>          | Click on <i>Chart output</i> to select.  |

Then click on *OK*.

You may find that your chart is slightly squashed, making the interval labels on the horizontal axis somewhat illegible. If so, click once on the chart, and you will see tiny black boxes emerge at each of the corners and mid-points of the chart border. Click and drag on the black box in the center of the lower border, stretching the

**IBM Monthly Returns: March 96 - January 2001**




**Figure 6.1a IBM returns histogram**

graph downward until the labels fit and can be read easily. You can produce a better title for this chart by clicking on the word “Histogram” and typing a more informative phrase, such as “IBM Monthly Returns: March 1996 – January 2001.”

Figure 6.1a is our version of the IBM histogram. Of course, there is no uniquely correct diagram; just aim for one that looks nice and conveys the sort of information you want for your readers.

**PG&E Monthly Returns: March 96 - January 2001**



**Figure 6.1b PG&E returns histogram**

Now do the same for another stock—say, PG&E—which is stored in Column N in your spreadsheet.

Figure 6.1b is a picture of what you might have ended up with.

*Warning: When you produce additional histograms, be careful to provide a different output range for each selection, like \$AA\$25 for the PG&E stock.*

Notice anything different about the two histograms?

## 6.5 Summary measures

Looking at the histograms in Figures 6.1a and 6.1b, you will notice some fairly systematic differences:

- (a) Returns for IBM are on average rather higher than PG&E, in the sense that the center of the distribution – however we define this – is more to the right. Call this the *central tendency* dimension.
- (b) The spread of observations is also wider for IBM than for PG&E or the S&P500 Gross index. Call this the *dispersion* dimension.
- (c) The histogram for PG&E is a bit crooked; this is the *skewness* aspect. In this example, it looks as though someone has taken a nice symmetric distribution and pushed it towards the right, thereby creating a long tail on the left. Hence, we call this *negative* skewness. If the long tail is to the right, the distribution is said to be *positively skewed*. Although it might be bit harder for you to pick it up simply by eyeballing, IBM is also a bit skewed, in this case positively.
- (d) The histogram for PG&E looks a bit peaky relative to a nice bell shape, with more weight in the center, less in the intermediate zones, and again more weight in the end zones. The pointy head / fat tails property is called kurtosis – specifically, *leptokurtic*. Rather like Dilbert’s boss.

Note that the bell-shaped property of the benchmark kurtosis comparison has a technical name - the *Normal* density. More on this shortly.

There are summary measures available for all of the tendencies mentioned above. In fact, there are several measures for each of the elements (a) – (d), but we’ll limit our discussion to the following:

- (a) To measure central tendency, we generally use the *arithmetic mean*: for a set of  $N$  numbers  $x_1, x_2, \dots, x_N$  we have:

$$\mu = \frac{1}{N} \sum_{i=1}^N x_i . \quad (1)$$

You will be used to the notation by now. Recall that the  $\Sigma$  is the summation sign: it means that you sum the observations, collectively denoted  $x_i$ , from observation number 1 to observation number  $n$ .

The *median* is another widely used measure of central tendency. It is the point where half the observations lie to the right (are greater than) and half to the left.

- (b) To measure dispersion, we generally use the *variance*, written as:

$$\sigma^2 = \frac{1}{N} [(x_1 - \mu)^2 + (x_2 - \mu)^2 + \dots + (x_N - \mu)^2] = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2 . \quad (2)$$

The square root of the variance, namely  $\sigma$ , is called the *standard deviation*.

*Note:* Often you will see the mean written as  $\bar{x}$  instead of  $\mu$ . Likewise, in the variance formula (2) you may see the divisor  $(N-1)$  used in place of  $N$ , while the sample mean  $\bar{x}$  is used in the formula in place of  $\mu$ . Indeed, this is the version used in the Excel functions; the result is often denoted  $s^2$  or  $S^2$ . These are called *sample means* and *variances*—they indicate that the data is being regarded as not a self-sufficient population, but a sample from some extended population of actual or potential data. Sample statistics (mean, variance) are used to estimate the population parameters. This is our first point of contact with the important topic of sampling estimates, and we'll get back to it in due course. For the practical work, go along with the Excel versions.

The *variance* is the “second order moment about the mean” – you can see that “second order” refers to the squares involved, and the “moment” is a measure of power or leverage. Notice that large deviations from the mean are squared up, or *powered up*, to become of much more moment in the scheme of things – hence, the variation aspect is increased by large deviations from the mean.

The *standard deviation* – sigma or  $\sigma$  – is another useful angle that we can call the “ballpark property.” It can be shown that at least 75% of the time, a random variable will lie within an interval equal to twice its standard deviation, and depending on the precise pattern of the data, sometimes a lot closer than that. So the standard deviation is a sort of bellwether as to how far the random variable  $x$  is likely to

depart from its mean. For most data patterns, you are not likely to find many observations further than one sigma from the mean, and very few more than two sigma's away.

(c) To measure skewness, we use the third order moment about the mean:

$$\gamma_3 = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^3 . \quad (3)$$

Again, the measures are powered up, in this case by the cube. But notice that positive and negative deviations from the mean are no longer treated the same – the cube of a negative deviation ( $x < \mu$ ) is negative, while the cube of a positive deviation ( $x > \mu$ ) is positive. So, if the histogram is nicely symmetric about the mean, there are as many positive as negative deviations, and they balance out to zero. On the other hand, a positive overall value for  $\gamma_3$  indicates that overall, there are bigger numbers above the mean than below it: the distribution is positively skewed. The same principle applies to negative skewness.

(d) To measure kurtosis, we proceed as follows:

First compute the fourth order moment about the mean:

$$\gamma_4 = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^4 .$$

Then compute the kurtosis measure as:

$$\text{kurt.} = \gamma_4 / \sigma^4 \quad (4)$$

Justifying this one is a bit less intuitive. It turns out for the ideal bell-shaped distribution (the *normal density*) that we used as a benchmark, the ratio in (4) turns out to have the value 3. Distributions that are peaky in the middle and fat in the tails (relative to the benchmark) have  $\text{kurt} > 3$ , while those that are fatter in the middle and skinnier in the tails have  $\text{kurt} < 3$ . Often the deviation ( $\text{kurt} - 3$ ) is in fact used and one simply observes whether it is positive or negative. This is essentially the case in Excel, where you just note whether the measure is positive or negative. If it's positive, the distribution is called leptokurtotic or *leptokurtic*; if it's negative, platykurtotic or *platykurtic*.

Once again, the issue of sampling arises. Sample data helps you *estimate* skewness and kurtosis in the population. In this case, you have to make some adjustments, typically involving divisors differing a bit from N. As the sample size gets larger, the difference from N gets smaller. As we noted before, Excel functions assume

the sample framework, and in the case of kurtosis, Excel uses the deviation from the normal benchmark of three to give a positive or negative outcome.

### Hands-on example

Try computing all the above measures for the stocks you have chosen, using Excel. First, let's compute the means.

(i) Computing the means.

Start by moving the pointer to an area of your worksheet well clear of your data and histograms; for instance, to cell [AM3]. Type the word **MEAN** in [AM3] and press **Enter**. Then calculate the mean (or arithmetic average) for each of the sets of data used in the histograms.

Excel has a statistical function which can do this for you. In cell [AP3] type **AVERAGE(C2:C60)** and press Enter. You have just computed the mean return for the IBM stock for the data period.

Alternatively, you can perform the same calculation using Excel's Function Wizard. Run the mouse pointer up to the standard toolbar and click on the *fx* button, which will generate a dialog box. Select "*Statistical*" from the *Function Category* list box, and "*Average*" as the *Function Name*. Then click on *OK*, type the range as **C2:C60**, and click on *Ok again*. (Note that instead of typing the range, you could have also selected those cells in your worksheet using the mouse or cell pointer.)

To produce the mean of a second series, you can simply copy the formula you have just created and paste it to the cell on the immediate right, using the cut and paste buttons in the standard toolbar. (Again, passing the mouse pointer directly over a toolbar button will reveal its name if you need a prompt.)

You could copy the formula into the next cell to the right for PG&E, but then you will have to edit the formula. Press the F2 key and change the cell range selection to **N2:N60** to obtain the PG&E data.

(ii) Computing dispersion, etc.

Excel has functions to compute the variance (VAR), skew (SKEW), and kurtosis (KURT), all of which can be called into action using the Function



Wizard. The desired function can be selected by clicking on the *fx* button, selecting *Statistical* as the function category, and selecting the required function name. Try this for the two selected stocks. You can incorporate all the measures into a table, with the variance, skew and kurtosis for a given stock appearing in the same column in the cells below the mean.

Finally, you can extend the work you've done by doing two more series. We suggest you use market returns and bill rate, to see whether they look different from the individual stocks. For the additional histograms, you can use the output range \$AA\$50 and \$AA\$65 for the two new selections.

To do the summary measures for the additional selections, you can simply copy all four measures for IBM and paste them across into the cells in the two columns directly to the *left* of the IBM calculations. (This copy and paste action will automatically result in the selection of the correct data—the data from Columns A and B of your spreadsheet.)

The result of applying the four Excel functions to the four variables is displayed in Table 6.1 below.

**Table 6.1 Summary comparison**

	CD rate	Market	IBM	PG&E
Mean	0.004662	0.015277	0.027725	-0.001301
Variance	1.6E-07	0.002161	0.010533	0.008008
Skewness	0.852361	-0.748774	0.340452	-0.714566
Kurtosis	0.169354	1.001306	-0.040204	2.123769

The results for IBM and PG&E conform to the conclusions we made earlier by eyeballing the histograms. For example, we see that PG&E is indeed negatively skewed, and IBM is positively skewed. PG&E is distinctly leptokurtic, while IBM is platykurtic (recall that Excel measures this from the zero base). You will notice some differences when you look at the market return and the CD rate. As you would expect, the dispersion of the market is less than for individual stocks, and the dispersion for the CD rate is quite a bit less.

## 6.6 Distribution function and densities

Now we are going to generalize things a bit. In essence, we will re-examine the data we have just explored, and view it as a sample from some underlying population of security returns. The underlying population can be provisionally thought of as a big histogram in the sky—but of course, this description would not satisfy our mathematician friends, so let's try to place our intuition on a firmer footing.


Go back to the histogram computation in the previous section. At each interval of the horizontal axis, instead of simply recording the number of observations falling within that interval, record the number falling less than the upper limit—that is, the number in the chosen interval plus the number in all the intervals less than this. Then express the result as a ratio, by dividing by the total number of observations (in this case,  $N = 59$ ) and plot. This is actually what you asked Excel to do when you selected *Cumulative Percentage* while producing the histograms for the two data series.

The result is called an *empirical distribution function*. The plot is cumulative in nature, since at each step, we are adding in the new histogram interval to all the preceding. We are in effect cumulating the histogram. The plot starts out at zero and ends at one, because by then all the numbers have been included. If you turn back to figures 6.1a and 6.1b, you can see that in each case, the empirical distribution function is listed as “cumulative %” and superimposed on the histogram.

This perspective is the key to making further generalizations. Suppose that instead of 59 observations, we have one thousand. You can see that we would have been able to subdivide our histogram intervals, and still have a reasonable number of observations in each. When you plot the empirical distribution function (i.e. the cumulative function), it will now look much smoother, because the steps are likely to be much smaller. You can see that if the number of observations gets very large, the curve would become very smooth. We can call this the *theoretical distribution function*. It is nice to think of the curve as continuous to emphasize the limiting process, but in fact it doesn't have to be continuous – it can have sudden jumps and still be a perfectly proper distribution function.

You might wonder what the data histogram would begin to look like. It turns out that the histogram as such does not survive the passage to infinity. The replacement model is a similarly shaped curve: bell-shaped, skewed, or whatever. You can interpret it to mean that the total area underneath the curve is one, and the area underneath the curve between any two points on the axis represents the probability that a single observation, selected at random, will fall within that interval. This is

called the *density function*. For technicians, the density function represents the mathematical derivative of the distribution function, or the distribution function represents the integral of the density function.



**Figure 6.2** Continuous distribution function and density

Figure 6.2 portrays the continuous distribution function,  $F(x)$ , and the associated density function,  $f(x)$ . The (cumulative) distribution function is now the nice smooth curve rising from left to right and reaching one. The corresponding density function represents the rate of change of the distribution function. The area under the density curve to the left of any point  $x$  is the probability of getting a value drawn at random that is less than  $x$ . This probability is given by the value  $F(x)$  of the distribution function.


Note that Excel has drawn the density function as a nice bell shape. The one illustrated above is in fact the Normal density, which is defined by the formula:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{1}{2\sigma^2} \cdot (x-\mu)^2},$$

where the number  $e$  is the base of natural logarithms.

The Normal density has just two parameters,  $m$  (the mean) and  $s$  (the standard deviation).

In Figure 6.3, Excel was used to explore the effect of changing  $\sigma$  and the area under the curve for different  $x$  values. Note that as  $\sigma$  increases from 0.5 to 1.0, the curve flattens close to the mean and the probability increases that an observation ( $x$ ) will fall somewhere in the tail—that is, further from the mean.



**Figure 6.3: Variational analysis on Normal density**

An especially important special case of the Normal density is where the mean is set equal to zero and the variance is set equal to one. This is called a *standard (or unit) Normal density*, and instead of  $x$  we conventionally use the symbol  $z$  for a standard normal variable. If you start with a variable  $x$  with arbitrary mean  $m$  and variance  $s^2$ , you can transform to a standard Normal by writing:

$$z = \frac{x - \mu}{\sigma}$$

In the case of a standard Normal density, the area under the curve to the right of  $z = 1.96$  is 2.5 %, and the area to the right of 1.96 and the left of -1.96 has a total of 5%. The point  $z = 1.96$  is called the 5% critical point for the Normal density.

Normal distributions and a closely related distribution called the *Student's t* are of great importance for testing hypotheses. Financial returns data usually aren't Normally distributed, but as we shall see, this does not stop us from using the Normal to test hypotheses about population parameters based on sample estimates.

With these concepts under our belts, we can reinterpret our data mining of section 6.4, referring to sample statistics which are used to estimate underlying population statistics (or *parameters*, as we will call them now). Once again, you will begin to see the divisor  $N-1$  used instead of the  $N$  used in variance formula (2), because we are now working with *sample* data. By replacing  $N$  with  $N-1$  in the sample calculation, we are correcting the  $s^2$  of the sample in order to estimate the underlying  $s^2$  of the *parent population*. A similar correction must be made to the denominators of the expressions for  $g_3$  and  $g_4$ . However, no such correction is necessary for the sample mean  $\bar{x}$  as an estimator of the population mean, which is usually denoted by the symbol  $m$ .

The description just given essentially motivates the theoretical distribution function and density in terms of a limiting process, as the number of observations in the sample gets large and larger. Actually, this represents a very particular view of probability known as the *frequentist interpretation*. While this view is a convenient illustrative device, it isn't the way most financial econometricians view the world. Econometricians generally have in mind a "data generation process" (DGP), which is a theoretical model couched at the outset in terms of probability densities. These models may never be able to generate nice stationary histograms such as we have discussed, so the passage to infinity, as we have described it, is impossible, and classical statistics do not necessarily apply.

Indeed, if you consider the context of our returns data, you can see why this is the case, at least roughly. Suppose that a stock market crash and subsequent backlash has occurred over the last several months. This is clearly a time of high volatility. Would your probability assessment of prices for the coming month be unaffected by this recent history? Clearly not. For one thing, you would certainly change (upwards!) the  $s^2$  for your distribution of the end of the month price. By doing so, you are starting to model the *conditional distribution* of the end of the month price, given the information available to you. In other words, you are modelling the conditional distribution of  $p$  at time  $t$ , given information available at time  $t-1$ . This is what econometricians are doing as well. In such a framework, the classical concept of a population out there somewhere is clearly inapplicable, since we really only ever have one observation possible for the coming month. The classical idea of a sample as 59 observations drawn from a parent population does have a meaning, but only in the sense of a stationary distribution that you might get by simply ignoring the temporal structure and considering what histogram boxes the prices fall into in the very long run (called the *stationary distribution*). The problem is, it does not always exist. We'll return to this later. For the moment, we shall continue to assume that classical statistics does apply and consider the important topic of sampling distributions.

## 6.7 Sampling distributions and hypothesis testing

Now suppose we assert at a staff meeting that the average overall monthly stock market returns (the S&P500, say) is  $m = 1\%$  and the standard deviation is  $\sigma^2 = 0.02$ . Note that we're asserting these as population parameters in general. Somebody challenges us; his name is Thomas Dubious. Tom is visiting from our Australian operation, and like many Aussies, he's a bit of a sceptic. He agrees with the  $s$ , but says that the correct theoretical mean  $m$  is  $0.7\%$ , much less attractive.

Let us take a closer look. Essentially, we have postulated an underlying (stationary) probability density for stock market returns with theoretical parameters  $m = 1\%$ ,  $s = 2\%$ . For the sake of argument, we will suppose that the density itself is nicely bell-shaped (Normal), in which case those two parameters are all that is needed to characterize the distribution. Thomas agrees with us on the shape and the variation, but not on the mean. How do we decide who is right?

The obvious thing to do is to look at some actual observations, which we shall take here as a random sample of stock market returns. Suppose we have been a bit lazy and taken a sample of just four observations. To estimate  $m$  we compute the sample average,  $\bar{x}$ . It comes out as  $\bar{x} = 1.1\%$ . Can we claim to have proven our point? Not yet. Our colleague Thomas would surely point out that our sample size was too small and that the average of just four observations could wobble all over the place.

Of course, he is right. The sample mean  $\bar{x}$  is itself a random variable, and it has its own theoretical mean and theoretical variance. The theoretical mean is exactly the same as that of the population from which you are sampling. But the *variance of the sample mean* is different – it depends on the size of the sample, as well as the variance of the population. As you would expect, the greater the sample size, the smaller the variation you would expect in the sample mean – it will wobble far less for a sample of size of one hundred than for a sample of size just four. Indeed, one can show that under pure random sampling, the theoretical variance of the sample mean is equal to the parent or population variance divided by the sample size (denoted  $n$  here):

$$\text{variance}(\bar{x}) = \sigma^2 / n \quad (5)$$

So, if  $s = 2\%$  and the sample size  $n = 4$ , the standard deviation of the mean (square root of the variance) is  $0.02/\sqrt{4} = 0.01$ , or  $1\%$ . You can see why our assertion that  $m = 1\%$  will carry little conviction; the true figure could be our doubter's  $0.7\%$  with an excellent chance as well.

Indeed, suppose Thomas right and the true value really is 0.7% (this condition is now our *null hypothesis*.) We can work out the probability of getting the observed sample mean of 1.1% given that  $m = 0.7\%$ . To do this, we use the fact that the sampling distribution of the sample mean is itself Normal. We work out the standard Normal value corresponding to  $x = 1.1$ , given a standard deviation of 1% and an assumed mean (the *null hypothesis*) of Tom's 0.7%. We get:

$$z = (1.1 - 0.7) / 1.0 = 0.4$$

Now you can go back to Excel, and using the Function Wizard, select "*Statistical*" as the *Function Category*, then "*NORMDIST*" as the *Function name*, and click on *Next*.

Then you must respond to the second dialog box as follows:

- |  |             |   |
|--|-------------|---|
| (1) on the first line next to <i>x</i> , type:     | <b>1.1</b>  | (then press the <b>Tab</b> key)                   |
| (2) on the second line next to <i>mean</i> , type: | <b>0.7</b>  | (press <b>Tab</b> )                               |
| (3) next to <i>standard_dev</i> , type:            | <b>1</b>    | (press <b>Tab</b> )                               |
| (4) next to <i>cumulative</i> , type:              | <b>true</b> | (press <b>Tab</b> ) and click<br>on <i>Finish</i> |

You will find that the function returns the value 0.655422, which represents the probability of getting an observed sample mean of up to the value 1.1. Hence, to find the probability of obtaining an outcome of 1.1 or greater, subtract the value returned by the function from 1.0 (i.e.:  $1.0 - 0.655422 = .344578$ ).

You find that the probability of getting an observed sample mean of 1.1 % or higher is 0.344578 or 34.6%. Thus, Thomas could very well have been right: even with his assumed mean of 0.7%, we still could have gotten a sample mean as high as 1.1% with an appreciable probability. In statistical terms, we could say that we have been unable to reject the null hypothesis that Tom is right, in favor of the alternative (that *we* are right).

On the other hand, with a sample size of  $n = 100$ , you can see from a re-application of formula (5) that the sample standard deviation of  $x$  is now  $0.02 / 10 = 0.2\%$ . Now calculate the  $z$  score for the sample mean of 1.1% with an assumed true mean of 0.7%. You will find that  $z = 20$ , and the probability of getting a  $z$  score as high as this is so small that it can safely be ignored. Almost certainly, Thomas is now wrong. Indeed, it turns out that even we have been a bit on the conservative side (to see this, assume a true mean of 1% and find the  $z$  score for the observed 1.1% on this basis).

Thomas, never one to give up, has a further objection. What if our ideas about the population standard deviation  $s$  – which you will remember we simply assumed was equal to 2% – were wrong, and in fact the true standard deviation is quite a lot higher? It is true that in this case the sampling variance of the estimator  $\bar{x}$  would be higher, and our conclusions would not be quite as devastating to his case. However, we can in fact estimate  $s^2$  by using the same data as we used to estimate the sample mean in section 6.2 using estimator (2). Remember, though, to use  $N - 1$  in the denominator since it is a sample estimate of a parent parameter.

A further complication is that the distribution of the sample mean is no longer Normal, and we have to use another set of tables (the *Student's t*). These are slightly more cumbersome, because the reference tables for the probabilities now depend on sample size. More specifically, they depend on the degrees of freedom in the data, which is equal to the sample size less one in this case. However, there are still 99 degrees of freedom and for a number this large, the Normal table is a very close approximation to the  $t$  table, so we may as well continue to use the Normal table. (Note that though we continue to use the term “table,” these days most of this data is accessible through Excel functions, making it unnecessary to reference an actual table.)

Let's proceed to estimate the sample standard deviation. We find it comes to 2.1. So the estimated sampling standard deviation of the mean  $\bar{x}$  is now  $0.21/10 = 2.1\%$ . Even with this change, you can verify that with the sample of size 100, not much will change as far as the  $z$  score – or the  $t$  score as we now call it – is concerned.

At this point, Thomas is out of objections, and we have won our case.

## Review

Although couched in the language of classical statistics, the discussion in this chapter has introduced some important ideas. Here are some of the topics we've covered.

1. The idea of an underlying population about which we are trying to make statements was introduced. In financial research, we typically replace the idea of a population, which is inherited from classical statistics, with the notion of an underlying model or *data generation process (DGP)*, the parameters of which we are trying to estimate or examine. Nevertheless, the language is much the same.



2. We also discussed the idea that we can make these statements by drawing some data from the population, or observing a more limited set of data presumed to have been generated by the underlying DGP.
3. A realization that once we use estimates based on samples, the values of those estimates can be expected to fluctuate depending on what numbers happen to appear in the samples – in other words, the sampling variation. This variation in sample-based estimates over different possible samples is called the *sampling distribution of the estimates*.
4. The sampling distribution of our estimates governs the degree of conviction with which we can infer things about the underlying parameters using the sampling estimates. Based on what we know about the sampling distribution of the estimates, we can make statements which give the probability that any hypotheses made on the basis of the sample data are correct.

Whew! We have covered quite a bit of the territory of statistical estimation and hypothesis testing, in a very intuitive way. For the purposes of illustration, we have used the sample mean as an estimator of a population mean. Later, we will deal with more useful sorts of parameters and estimators with respect to financial data. But the fundamental concepts are the same, no matter what the context. In the meantime, note that all the standard estimators you are likely to use come with estimates of their sampling distributions attached. In particular, the package you use will almost certainly give not only the estimate itself, but also an estimate of the its sampling variance. These should always be presented together, so that the reader can first look at the estimated value and then at the sampling variance of the estimate to see how likely it is to wobble around. This allows the reader to make up his or her own mind as to how much credence to give the estimate.

This is all the formal theory of sampling distributions that we need. We've also touched on the topic of hypothesis testing in the context of our debate with Thomas Dubious. However, we shall leave a formal treatment of hypothesis testing to the statistics textbooks.

1. Not too long ago, we saw an MBA questionnaire circulated to corporate treasurers asking them whether they practiced formal risk management techniques. Would the following categories of recipients have the same response rate?
  - (a) Treasurers who do not practice risk management techniques
  - (b) Those who do practice risk management techniques, but are too busy and/or do not like MBA students
  - (c) Those who are not too busy and who do practice risk management techniques

Hence, evaluate such a questionnaire as a research tool.

2. Although we have followed convention in using the arithmetic mean of returns, for returns data, the geometric mean is often – and more correctly – used. The geometric mean  $g$  is defined implicitly by the formula:

$$(1 + g)^N = (1 + x_1)(1 + x_2) \dots (1 + x_N)$$

In other words, it is the constant return that would compound up to exactly the same terminal amount as the observed sequence of returns  $x_1, x_2, \dots, x_N$ .

Comparative returns data are used extensively by the research firms that compile fund rankings. For such purposes, which do you think is the more correct measure—the arithmetic mean or the geometric mean?

3. Here are some random variables of different kinds. What sort of density shape do you think each would have? (Look in particular at the symmetry, kurtosis angles, and whether the shape might be *bimodal*, meaning it might have two peaks rather than one.)
  - (i) Daily stock market returns.
  - (ii) The weights (in pounds or kilos) of the population at large.
  - (iii) The weights of money market dealers.
  - (iv) Incomes in the population at large.
  - (v) The number of tourists eaten by crocodiles in Queensland each year.
  - (vi) The distribution of grades in your class on the next exam.

4. We know from the Normal tables that for a standard Normal density,  $\text{prob}(z > 1.96) = 0.025$ . Because the Normal density is symmetric, we can also say that  $\text{prob}(z < -1.96) = 0.025$ . So  $+1.96$  and  $-1.96$  are the *upper and lower 2.5 % critical points*, making up 5 % together, so that they are sometimes called the two-tailed 5 % points.

Assume that the distribution of a sample mean  $\bar{x}$  is Normal with mean  $m$  and variance  $\sigma^2/n$ . Use the standard Normal limits above to claim:

$$\left[ \bar{x} - 1.96 \frac{\sigma}{\sqrt{n}}, \bar{x} + 1.96 \frac{\sigma}{\sqrt{n}} \right]$$

Suppose you could repeatedly draw samples of size  $n$  and compute the mean of each such sample. Then 95% of the time, the interval will contain the true mean  $m$ . The above interval is said to be a *95% confidence interval for  $m$* . Loosely speaking, once we have computed the sample mean  $\bar{x}$  we can be 95% confident that the true mean will lie between the two limits.

Getting back to our debate with Thomas, let's look at this another way. We draw a sample of size four, and find a sample mean of 1.1%. We know the population standard deviation is  $\sigma = 2\%$ .

- (i) Find a 95 % confidence limits for the population mean of stock returns.
- (ii) Now change the sample size to  $n = 100$ . What happens to the confidence limits?



---

# Seven

## Regression and All That

In this chapter, we extend our discussion to consider the association between two or more variables, rather than each one in isolation. It is important for us to do so, because many of the statements in financial theory and practice are indeed about relationships. Often these appear in the form of hypothesized equilibrium properties. For instance, the famous CAPM theory concerns the relationship in a state of capital market equilibrium between the return on individual stocks on the one hand, and the market return, on the other. Or, when we hedge one position with another, we need a formal way of handling that vital *stochastic* (a fancy name for random) relationship between the hedge instrument and what is being hedged. Regression methodology is a window into all this.

## 7.1 Notation and definitions

Covariance	The average product of deviations from means of two variables; measures the extent to which one variable moves along with the other (positive covariance) or in opposition to the other (negative covariance).
Correlation coefficient	The coefficient of correlation between two variables; equal to the covariance divided by the product of the standard deviations, with value lying between +1 and -1. If the two variables have no statistical association, the value is zero, while if the two variables plot perfectly along the straight line, the value is either plus one or minus one.
Degrees of freedom	The effective independent dimensionality of a computed statistic, which may be less than the number of observations.
Histogram	A bar graph of a frequency table, where the intervals of the table plot into bars or boxes, with a base equal to the interval and height equal to the frequency.
Null hypothesis	The value of a parameter given an assumed base case model or condition, to be accepted or rejected for the alternative hypothesis—its value under an alternative condition.
OLS	Acronym for Ordinary Least Squares, the most basic statistical technique for fitting parameter values to a line or curve; operates by choosing the parameter values that minimize the sum of squared deviations from the assumed line or curve.
Regression	The hypothesized relationship between a dependent or left hand variable, and one or more independent, or right hand variables, plus a residual disturbance term to allow for random influences.
Scatter plot or diagram	A graph of data pairs with observations of one variable along the horizontal axis and observations on the other along the vertical axis.

## 7.2 Bivariate data exploration

Now take two of the disk series, say an individual stock of your choice (call it  $y$ ) and the market (call it  $x$ ). We are going to explore some measures of association between them. As in classical statistics, it is convenient to imagine that the 59 observations form a sample from some stationary bivariate distribution, so we shall use the sample versions of the various estimators, which you will recall typically have  $N-1$  rather than  $N$  in the denominators, at least for second order variational measures.

Before we go further, a bit of data exploration in the form of a scatter diagram may be helpful.

Copy your data on the “Market Returns” from column B of your open Excel worksheet into another free zone on your worksheet, hitting the paste button when you have run the cell pointer across to, say, cell [BA2]. Copy data from column C (IBM Ltd), pasting it into cell [BB2] on down, so that it lies directly to the right of the market returns data. In cell [BA1], type the word **Market**; in cell [BB1], type **IBM**. Employing the Excel Chart Wizard, we will incorporate this newly isolated data into a scatter diagram, which has the market on the horizontal and the IBM stock on the vertical. The individual points will appear as small diamond shapes.

### Example 7.1

Move the cell pointer to the point where you want your graph, for instance cell [BC2], and with the mouse pointer click on the *Chart Wizard* button in the standard toolbar to bring up the first of five dialog boxes. (Alternatively, from the menu bar, select *Insert, Chart...*, and proceed as follows.)

Step 1: Click on the XY Scatter Diagram to select the *Chart Type*. Then select the first format for the XY Scatter Chart (which has no connecting lines between the data points), and click on [Next>].

Step 2: Click on cell [\$BA\$1] and holding down the left button of the mouse, drag to cell [\$BB\$60] to define the data range for your graph. (Or simply type the range with a colon between the start and end points, but note that the range should be specified with absolute addresses.) Check that the “Data series in:” shows as *Columns*. Click on [Next>].

Step 3: A sample chart is displayed on the screen, along with the “Chart Options” dialog box. Under the *Titles* tab, for *Chart Title:*, type in a suitable heading such as: **IBM Ltd vs. Market Return**.

For Value (X) Axis, type: **Market Return**. For the Value [Y] Axis, type: **IBM**.

Press the “Legend” tab, and click on the *Show legend* tick box to *remove* the tick (thereby removing the default legend). Click on [Next>].

Step 4: For chart location, select to place the chart *as object in* the current worksheet. Then click [Finish].

The scatter chart appears, spreading to the east and south of cell [BC1]. It may appear rather squashed, but this can be adjusted by clicking on the appropriate border and dragging the small black box to the desired location. You may find it necessary to drag the lower border down a number of lines, making the chart look tall and skinny. Then, to give the chart a slightly more normal-looking shape, customize the vertical axis as described below.

### Customized axis

Left click on the Y-axis to select it, making sure that Excel has indicated its selection with the message “Value (Y) axis.” Then right click with the Y-axis selected to call up the *Format Axis* dialog box. Choose the *Scale* options. Adjust the minimum to -0.25 and the maximum to 0.35, and click [OK]. Equalizing the horizontal and vertical scale of the diagram (so that a measure of 0.05 on the horizontal axis is the same distance as a measure of 0.05 on the vertical axis) will make your graph easier to read. This can be achieved using click and drag adjustments on the borders. Figure 7.1a is the resulting scatter diagram.






Figure 7.1a Scatter diagram of IBM against market returns

### Example 7.2

You can do a similar scatter chart for another stock—say, Wendy’s—which resides in column P of your worksheet. Copy the data for the market into cell [BA64] and down, and the data in column P for Wendy’s in cell [BB64] on down. In cell [BA63], type the word **Market**; in cell [BB63], type **Wendys**. Run the mouse pointer to cell [BC63] and click to anchor a new scatter chart. Then proceed as above using the newly positioned data.

Start by clicking on the Chart Wizard button and specifying the range \$BA\$63:\$BB\$122. The title for the new chart could be “Wendy’s vs. Market Returns,” with “Wendy’s” on the y-axis and “Market Returns” on the x-axis. Again, you will probably find it helpful to customize the y-axis, this time by specifying a minimum of (-0.2) and a maximum of 0.3. You will probably also wish to re-size the graph by dragging the borders as with the previous scatter diagram. Figure 7.1b reproduces a scatter diagram for Wendy’s.

#### Modify and beautify

You might want to modify the chart title or axis labels at some time after the graph is constructed. For example, suppose you wish to alter the title in the second chart above to “Scatter Diagram: Wendy’s International vs. Market.” Click (using the left mouse button) on the original title to highlight it. Then click within the title after it is highlighted and proceed to edit.

Other aspects of the graph presentation, such as borders and shading, may also be changed through formatting. There are two main format sections to your graph: the inner plot area which contains the graph itself, and the outer chart area which holds the title and axis labels. To change either of these zones, first click anywhere on the chart to select it.

If you wish to alter the plot area, left click within the plot area to highlight it (Excel will signal with the indicator tag reading “Plot Area”) and then double click the left mouse button to activate the *Format Plot Area* dialog box. (Rather than double clicking, you can right click and select *Format*

*Plot Area...* instead.) You can then adjust either the border or the background pattern of the plot area via the dialog box. To change the main chart area, left click in that section to highlight, and then double click to activate the *Format Chart Area* dialog box. That takes care of appearances.

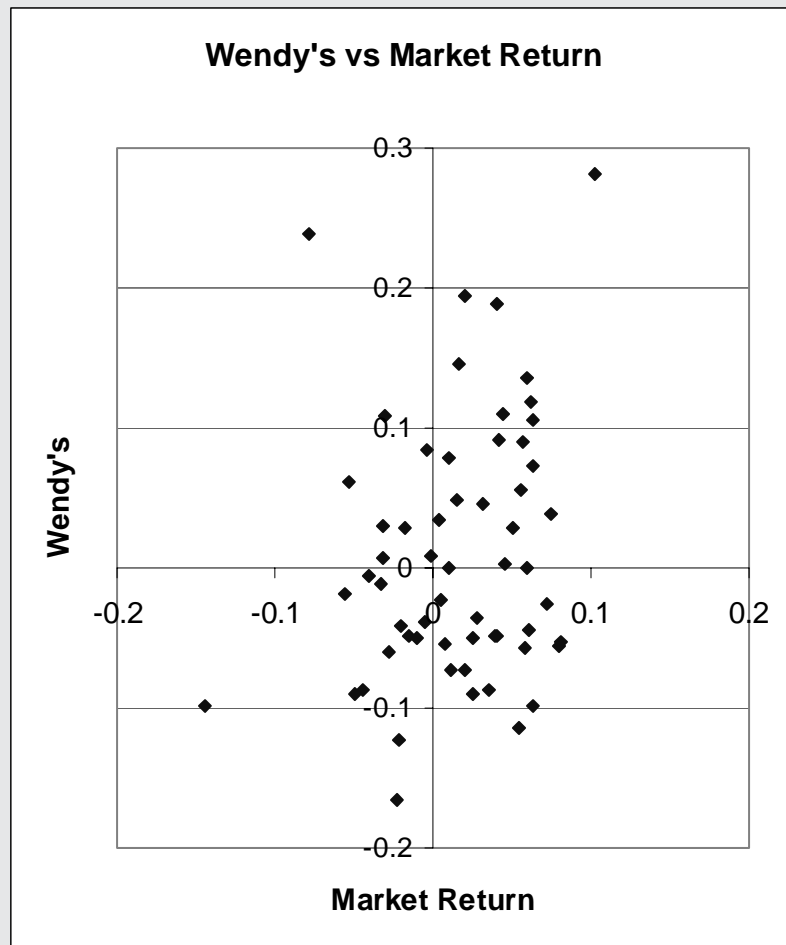


Figure 7.1b Scatter diagram of Wendy's against market returns

### 7.3 Regression statistics

Now think about the following statistic:

$$\sigma_{yx} = \frac{1}{N-1} \sum_{i=1}^N (y_i - \bar{y})(x_i - \bar{x}) .$$

This is called the covariance between  $y$  and  $x$ ; it is symmetric so that we could have written it as  $\sigma_{xy}$ . Use of the  $N-1$  rather than the  $N$  in the denominator means that we are looking upon it as a sample estimate of a parent  $\sigma_{yx}$ , so we really should write it as  $S_{yx}$  or something similar. However, the precise denominator turns out not to matter in the present context.

The covariance is the sum of the deviations from the mean for each variable, multiplied together. Suppose that positive deviations for  $y$  tend to be associated with positive deviations in  $x$  – in other words, when  $x$  is higher than usual, so is  $y$ . Then the covariance will be positive. But if positive deviations in  $x$  tend to be associated with negative deviations in  $y$ , then  $\sigma_{yx} < 0$ .

By itself, the covariance is not used much as an end product, but it is an essential intermediate step for a number of statistics that are. Two such are as follows:

1. The correlation coefficient:

$$r = \frac{\sigma_{yz}}{\sigma_y \sigma_x} .$$

Sometimes you will find this written as the uppercase  $R$  or as its square  $R^2$  (See Section 7.6).

2. The regression coefficient of  $y$  on  $x$ :

$$b = \frac{\sigma_{yx}}{\sigma_x^2} = \frac{\sum_{i=1}^N (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^N (x_i - \bar{x})^2} .$$

The correlation coefficient is a measure of the closeness of association between  $y$  and  $x$ . More precisely, it measures the degree to which  $y$  can be explained as a linear function of  $x$ . Thus we could write something like the following:

For each observation  $i$ ,

$$y_i = \alpha + \beta x_i + \varepsilon_i . \quad (4)$$

Here,  $y_i$  would be explained in terms of a linear function in  $x_i$ , which is the first part of the right hand side. The parameter  $\alpha$  is called the *intercept*, the point where the line crosses the vertical axis. The parameter  $\beta$  is called the *slope*, which measures the steepness of the line. For example,  $\beta = 1$  would mean that  $y$  changes one for one with  $x$ . Beta could be negative, in which case the line would slope downwards. Figure 7.2 illustrates this example. The last component on the right hand side of (4) (the residual  $\varepsilon_i$ ) simply represents the idea that the proposed line does not fit exactly. If it did, for all intents and purposes,  $y$  would just be a rescaled version of  $x$ —hardly a separate variable in its own right.

If the correlation coefficient were one, then the line would indeed fit exactly. If the correlation coefficient is zero, then the line – any line – is simply no help at all. If in the latter case we plot the actual observations of  $y_i$  against  $x_i$  as a scatter diagram, then the pattern is likely to be just a blob, with no structure. Actually, there could be a non-linear structure to the relationship, but because the correlation coefficient is just a measure of linear associations, it would not pick this up.

The regression coefficient  $b$  is best regarded as a sampling estimate of the slope coefficient  $\beta$ . Technically, it is a special sort of estimate called the least squares estimate. This means that to fit (4) to an observed scatter diagram, you could think of it as trying various values of the  $\beta$ 's and juggling them around to end up with the values that minimize the squared sum of fitted residuals from the line (i.e.  $\sum e_i^2$ ). As you can see in Figure 7.2, line A (i.e.  $\alpha^1, \beta^1$ ) is a much better fit than line B ( $\alpha^2, \beta^2$ ). The best fit is given by  $\beta = b$ , the regression coefficient, with the estimate for  $\alpha$  defined by:

$$a = \bar{y} - b\bar{x} .$$

Incidentally, this relationship indicates that the line of best fit must pass through the sample mean point on the scatter diagram which is the point:

$$(\bar{y}, \bar{x}) .$$

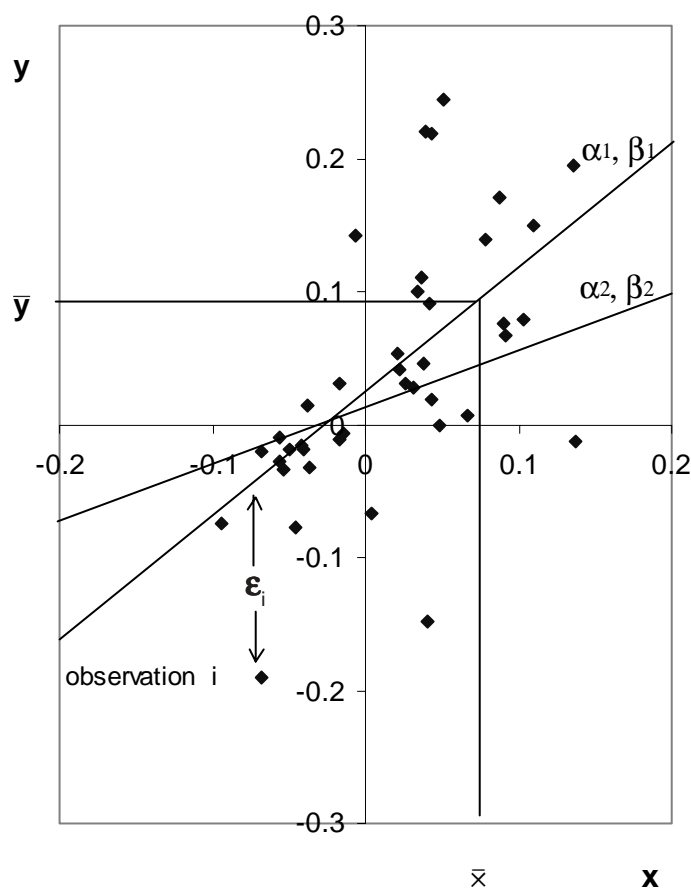


Figure 7.2 Regression line on a scatter diagram

### Example 7.3

Can we do all of this with Excel? Of course! Let us continue with the same data series that we used in Examples 7.1 and 7.2, namely IBM and Wendy's. In each case, we will employ the market return as the associated variable for the purposes of regression and correlation examples. Actually, we will see in Section 7.4 that as a financial procedure, this is not 100% sensible, but we'll get to that in due course; for the moment our purpose is purely illustrative.

The correlation coefficient tool is found in the menu toolbar by clicking on *Tools, Data Analysis...*, and then selecting *Correlation* from the "Analysis

tools” box, and clicking [OK]. (If *Data Analysis...* does not show up, you will need to select *Tools, Add-Ins*, and tick *Analysis ToolPak*, and click on *OK* before proceeding with Example 7.3.)

This action will generate a Correlation dialog box for which you can nominate the Input range [\$BA\$1:\$BB\$60], to use IBM and Market returns.

The input data is grouped by:

- Columns, with
- Labels in First Row.

For output select:

- Output Range (as) \$BW\$2 and hit [OK].

To do the regression analysis, slide the cell pointer over to cell [CA2] and in the menu toolbar click on *Tools, Data Analysis...*, and then select *Regression* from the “Analysis tools” box, and click on [OK].

To continue the analysis using IBM, nominate [\$BB\$1:\$BB\$60] as the Input for the Y Range, and [\$BA\$1:\$BA\$60] as Input for the X Range. Then select the following:

- Labels

For Output Options, select:

- Output Range \$CA\$2

For Residuals, select:


- Residuals
- Line Fit Plots

Then click [OK].

The above regression analysis procedure will generate a set of statistics and another scatter diagram with a fitted regression line. For the moment,

we will simply bypass commentary on the set of statistics that flashes up – we will deal with these in the extended example of the next section.


Again, it may be necessary to tidy up the diagram by dragging on borders and adjusting the scale on either axis. You might also wish to alter the look of the data points. To do so, left click on the plot area of the graph, and then double click on a data point. This action will open up the *Format Data Point* dialog box. Select *Patterns* to change the look of the set of data points. If you have clicked on the predicted points, you can change the shape, color and so forth of the fitted line. Or, by double clicking on one of the “IBM” points, you can use the resulting dialog box to alter the presentation of these points.



**Figure 7.3a** Fitted regression line for IBM on market returns



Now try the above procedure on the Wendy's data paired with the Market Return. The resulting Excel graphs for IBM and Wendy's are displayed as Figures 7.3a and 7.3b.



**Figure 7.3b Fitted regression line, Wendy's on market returns**

#### 7.4 More regression theory: goodness of fit

Referring back to the representation (4) of the regression model, one can estimate the residual terms  $e_i$  corresponding to each observed  $x_i, y_i$  as:

$$e_i = y_i - a - b x_i$$

where  $a, b$  are the least squares estimates.

The sum of squares  $SSE = \sum e_i^2$  is called the error sum of squares, and the quantity

$$\hat{\sigma}_\varepsilon^2 = \frac{\sum e_i^2}{(N-2)} = \frac{SSE}{(N-2)} \quad (5)$$

is described as the *fitted sigma squared*. We subtract two (2) from N in the denominator of (5), because we had to estimate two parameters, namely  $\alpha$  and  $\beta$ , before we could calculate the statistic we are after. In statistics terminology, the computation uses up two *degrees of freedom*. In Excel, the fitted sigma squared is referred to as the *standard error*, that is, the standard error of the equation as a whole.

One can show that:

$$\hat{\sigma}_\varepsilon^2 = (1-r^2)\sigma_y^2$$

This means that the estimated error variance is just that of the dependent variable, reduced by the factor  $(1-r^2)$ . You will notice that if  $r = 1$ , the sum of squares of the residuals is just zero, which can only be true if each of the  $\varepsilon_i$  is zero. The fitted line is in fact a perfect fit, reinforcing an earlier remark to that effect.

The regression parameter of primary interest is usually the slope parameter  $\beta$ . The OLS (*ordinary least squares*, which we have just calculated) estimator for  $\beta$  is the regression coefficient  $b$ , computed automatically by the program. Because  $b$  is a sample estimator, it has its own sampling distribution. If you knew the residual variance  $\sigma$ , you could compute the  $z$  score for your fitted value of the regression coefficient. Usually, you don't know  $\sigma$  and it has to be estimated by the program; in this case, the  $z$  score is replaced by the  $t$  score and the probability tables are those of the Student's  $t$  distribution. However, the statistic is formally the same.

A common null hypothesis is that the true slope  $\beta$  is zero – which is a way of saying that the chosen right hand variable  $x$  (the *explanatory* variable) has no effect on the dependent variable  $y$ . In this case, the  $t$  score is simply the observed regression coefficient  $b$  divided by the estimated sample standard deviation of  $b$ ; the latter is printed out by the program. The resulting ratio is often loosely referred to as “the  $t$  value.” You may recall from chapter 6 that the  $t$  distribution and tables are dependent on the degrees of freedom. In the present context, these are taken as the sample size less the number of primary regression parameters, of which there are two here,  $\alpha$  and  $\beta$ . However, if the number of degrees of freedom is greater than 50 or so, the  $z$  score and Normal table will do very well. A good rule

of thumb is that if the computed  $t$  score as produced by the program is greater than 2.00, it is interpreted as rejecting a null hypothesis of “no effect” in favor of one where there *is* an effect—that is,  $\beta \neq 0$ .

You can try this process on our extended Examples 7.1 – 7.3 above. We shall not do it here, because we will be showing it in a more interesting framework in the next section.

## 7.5 The CAPM beta

Now we are going to use the information above to do something useful—namely, to fit CAPM betas for the stocks in the data file. If you are not familiar with the idea of the CAPM beta, you could refer to most texts on investments, which would provide extended coverage. For our purposes we will assume that you have some knowledge in this area. (Though the empirical treatment in many elementary texts leaves something to be desired.)

First, some comments on notation. Remember that we have 14 stocks, so we will now use the index  $i$  to refer to the representative stock. Since we are dealing with observations in the form of time series, we will reserve the index  $t$  for the representative observation (earlier we used  $i$  for this purpose, but here it is more natural to reserve this for the individual stock). Our notation will be as follows:

$r_{i,t}$  = return on stock  $i$  for month  $t$

$R_t$  = return on the market for time  $t$

$r_t$  = risk-free rate, which we shall take here as the one-month CD rate

Note that the certificate of deposit (CD) rate is not the ideal thing to use for the risk free rate; even short dated CDs are not technically “free” either of credit risk or of interest rate risk. However, the existence of deposit insurance reduces credit risk, so we will stick with what is available.

The CAPM, in the form we need it, takes the form of two equations:

$$r_{i,t} = \alpha_{i,t} + \beta_i R_t + \varepsilon_{it} \quad (6a)$$

for each stock  $i$ , and:

$$\mu_{i,t} = \rho_t + \beta_i (\mu_{R,t} - \rho_t) \quad (6b)$$

Equation (6a) assumes that for each stock  $i$ , there is a linear relationship between the stock and the market. The parameter  $\beta$  is the theoretical regression (slope) coefficient, and this is supposed to be constant over the period of the sample – you will note that it does not depend on the time index  $t$ . On the other hand, the intercept term in (6a) cannot be independent of time. We shall see why in a moment.

Equation (6b) is the core of the CAPM hypothesis. It says that if there is a linear association such as the one represented in the first equation, then the stocks must trade so that their expected return  $\mu_{i,t} = E r_t$  at the start of period  $t$  bears a particular relationship to the expected return on the market  $\mu_{R,t} = E R_t$ , again as of time  $t$ . The expectations are dependent on information assumed to be available at the beginning of period  $t$  to investors. Note that this will include the value of  $\rho_t$  which is by definition known as of that instant – otherwise it would hardly be a risk-free rate. All this emphasizes that the distributions in which the CAPM is couched are essentially *conditional* probability distributions – conditional, that is, on information presumed to be available to the investor at the start of time  $t$ . (The notion of conditionality is highly important in finance and we will come back to it in the next chapter).

According to the CAPM, the beta from the regression, which is Equation (6a), gets transferred through to the conditional expectation of the return, which is Equation (6b). If we take expectations of both sides of the regression relationship (6a), the conditional expectation of the residual term is zero, so we have:

$$\mu_{i,t} = \alpha_{i,t} + \beta_i \mu_{R,t} \quad (6c)$$

Now you can use (6b) and (6c) to find that:

$$\alpha_{i,t} = (1 - \beta_i)\rho_t \quad (6d)$$

You can see why we could not specify that the parameters of the regression (6a) could both be constant, independent of time. It's OK for the beta but not for the alpha. This is also why you cannot estimate  $\beta_i$  as the simple regression coefficient of the stock on the market, which is a mistake that is often made. Notice, however, that if you substitute (6d) in (6a), you end up with:

$$r_{i,t} - \rho_t = \beta_i (R_t - \rho_t) \quad (7)$$

This means that you can estimate the betas by regressing the *excess returns*  $r_{i,t} - \rho_t$  on the market excess return  $R_t - \rho_t$ . However, you should do so without an intercept term. Equation (7) is the *time series equation of the CAPM*.

**Example 7.4**

So let's put this into action. We have picked General Motors (GM) as our first stock. First we have to create new variables by subtracting off the risk-free rate (one-month CD rate) from both the GM return and the market.

For this section, it will prove helpful to create a new worksheet to hold data, and new calculations and graphs. Run the mouse pointer up to the menu toolbar and select *Insert, Worksheet*. To clearly label the new worksheet so that its contents are obvious at a glance, it could be given a more descriptive name than "Sheet 2." From the menu toolbar, select *Format, Sheet, Rename...* and type, say, **BETAS** into the "Name" box. Then click on [OK]. It would also make things clearer at this point to rename the original Sheet 1 with a suitable title such as "RETURNS." Click on the "Sheet 1" tab at the bottom left of the work file to return to the original worksheet. Then follow the above procedure (*Format, Sheet, Rename...*) to give this worksheet the name **RETURNS**.

Our first task will be to subtract the risk-free rate for the appropriate time period from each of the market returns and stock observations, and bring this adjusted data onto the BETAS worksheet. Here is an easy way to do so.

Starting in cell [A2] of the BETAS sheet, type the character = and then click on the RETURNS worksheet tab which appears at the bottom left of the window (but don't hit [Enter] at this stage). Now, in the RETURNS worksheet, run the cell pointer to cell [B2], type - (the minus sign), click on [A2], and hit [Enter].

These steps will bring you back to cell [A2] of the BETAS sheet. In this cell, you will now see the result of the subtraction operation you have constructed: 0.030745. In the formula bar, you will observe the expression "=RETURNS!B2-RETURNS!A2." Run the mouse up to the formula bar, and click to the left of the "A2." Then type a \$ in front of the "A2" to make the column part of the address absolute (but *not* the row part). Hit [Enter]. Next, copy the formula in this cell down to all the cells in the column as far as [A60]. You have just created a set of data which shows the market return less the "risk-free rate" for 59 observations over time. Now run the cell pointer up to cell [A1], and type the label: **Market**.

Using the mouse pointer, with a click and drag action, highlight the entire set of 59 observations in column A. With the left mouse button held down, move the mouse to the very bottom right hand corner of the data so that it shows as a small black cross and drag across to column P. You will now have a full set of market returns, individual stock returns, and bond returns from which the appropriate CD rate has been subtracted. These represent our approximation of the excess return over the “risk-free rate” for each observation.

The scatter diagram of the excess returns on General Motors versus those of the market is shown as the black diamonds of Figure 7.4. Clearly there is a line of best fit, and its slope looks to be roughly one.

Now we follow the procedures for regression analysis. Slide the cell pointer over to cell [A1] and in the menu toolbar click on *Tools, Data Analysis...*, and then select *Regression* from the “Analysis tools” box, and click on [OK].

To continue the analysis using General Motors, nominate [C\$1:C\$60] as the Input for the Y Range, and [A\$1:A\$60] as Input for the X Range. Then select the following:

Labels

For Output Options, select:

Output Range (and specify an area clear of the data, such as T\$1).

For Residuals, select:

Residuals       Line Fit Plots

Then click [OK].

You will find a great deal of output, some of which has been reproduced on the following page as Table 7.1

**Table 7.1: Selected regression output, General Motors and CAPM**

## SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.525975316
R Square	0.276650033
Adjusted R Square	0.263959683
Standard Error	0.08273044
Observations	59

## ANOVA

	df	SS	MS	F	Significance F
Regression	1	0.149206514	0.1492065	21.800031	1.88381E-05
Residual	57	0.390126564	0.0068443		
Total	58	0.539333078			

	Coefficients	Standard Error	t Stat	P-value	Lower 95%	Upper 95%
Intercept	-0.00424123	0.01105167	-0.3837637	0.7025814	-0.02637183	0.0178894
Market	1.089388768	0.233321272	4.6690504	1.884E-05	0.62217078	1.5566068

## RESIDUAL OUTPUT

Observation	Predicted GM	Residuals
1	0.029251947	0.020112501
2	-0.004755	0.028288839
3	-0.09636927	-0.104268423
4	-0.0145354	-0.035153698
5	-0.06746167	-0.009433566
6	0.056926898	0.175803676
7	-0.02700902	0.002158509
8	0.016865913	-0.20023494
9	-0.03294481	-0.212783954
10	-0.04229266	0.167768536
11	0.101770777	-0.018036883
12	-0.03481739	-0.01969153
13	-0.06383463	0.167391052
14	0.054088916	-0.07439996
15	0.012629996	0.037572439
16	0.060184641	0.054513584
17	-0.03904002	-0.015402437

Among the useful elements in Table 7.1, at least for our current purposes are the following:

1. The fitted beta has value 1.089. Its standard error is 0.233 or something close to that. The *t value* is 4.669; you will recall that this is really the *t score* for testing the null hypothesis of no effect (i.e.  $\beta = 0$ ), and the rule of thumb that a *t value* greater than 2.0 in absolute value means we reject the null hypothesis. So the beta is clearly significantly different from zero.

You could form the  $t$  score against a null hypothesis that the true value of beta is 1.0. The  $t$  score would then be:

$$t = (1.089 - 1.0) / 0.233 = 0.38192$$

Since this is less than the benchmark critical point of 2.00 for 57 degrees of freedom, we can accept the hypothesis that the true beta is equal to one. If you formulate the test as a so-called *one-tailed* test, where the only allowable alternative hypothesis is that beta is greater than one, then the corresponding critical value for the  $t$  distribution reduces to 1.67, rather than 2.00. Because our  $t$  value is less than this, we would reject the alternative that beta was greater than one. Things become a little less distinct in cases where there are marginal values.

2. The intercept – the value of alpha – is -0.004. You can test the null hypothesis that it is equal to zero – as the CAPM predicts – by forming its own  $t$  score. This is done automatically by the program to give a  $t$  value of -0.384, which is not significantly different from zero. Thus this fits, so to speak, with the CAPM.
3. The  $r$  statistic is printed out as “multiple R.” We’ll explain what this means in the next section, but here it can be taken as the  $r$  value, and is 0.526. The corresponding  $r^2 = 0.277$ , which means that 28% of the variation in General Motors’ excess returns is explained by the market.
4. The estimated value of the standard error (which appears as the square root of the variance) is  $\sigma_\epsilon = 0.083$ .

Figure 7.4 below plots the scatter diagram of the excess returns, with General Motors on the vertical axis and the market on the horizontal. The fitted line is also shown.

Don’t bother too much about the rest of the output. Anova is a measure of whether the relationship as a whole is statistically significant. In the present case, the only right hand variable is the market return, and its  $t$  value will do for our purposes. The program also prints the fitted residuals. It is a good idea to plot them against time to see whether there is some temporal pattern. The existence of such a pattern might mitigate the view that the residuals forming the fitted relationship are purely random. In more advanced treatments, there are formal tests for the existence of temporal relationship



among the residuals, such as the Durbin Watson test, but this is a little beyond our scope here.

Overall, it appears that the CAPM does well, at least so far as GM is concerned. The value of beta as a bit greater than one indicates a slightly aggressive stock. This is more or less as we might expect, because GM is a car producer and might be expected to vary procyclically with the general business cycle and through this, with the market return index.




Figure 7.4 CAPM graphics – GM against the market

Try some or all of the other stocks. What about Praxair or Wendy's, for instance? Can you rationalize the betas you get for these companies?

One final point: If the CAPM relationship specifies no intercept – and Equation (7) does not have one – then you should fit the equation by suppressing the intercept. Of course, we did have an intercept in the GM case, which is standard, but it came out as statistically insignificant nonetheless, just as we would hope. However, it is not sufficient to regard the fitted slope parameter from such an equation as the beta. We really have to fit the whole equation without the intercept, and then look at the fitted slope parameter. This can be done as follows.

Follow the usual procedure for performing regression analysis on Excel, as described in Examples 7.3 and 7.4, with one small difference. Having selected *regression* from the *Analysis tools* and specified the Y Range and the X Range, select:

✓ Labels, as usual, but in addition, select:

✓ **Constant is zero.**

Continue with your standard selection of the output options and residuals, and then click [OK]. Your resulting regression line will be forced through the origin. If you do this for the GM excess returns data against the market, your output will change by a small amount. Try it and see.

## 7.6 Regression extensions: multiple regression

Suppose we thought that variations in stock returns might be explicable in terms not only of the stock market as a whole, but also of influences from outside the stock market. In financial general equilibrium jargon, we could say that the stock market did not *span* the systematic sources of influence on stock returns. For instance, we might decide that there are two factors that affecting stock returns. One might be a general factor just for stocks. The other might be a factor that comes over from another market, such as the bond market. Or it could be an exchange rate factor, representing external influences on our chosen stock or set of stocks. Since we have data for the bond returns, we can investigate this.

We might set up the following extended version of the CAPM equation:

$$r_{i,t} - \rho_t = \beta_i (R_t - \rho_t) + \gamma_i (r_{b,t} - \rho_t) + \varepsilon_{i,t} \quad (8)$$

Here, everything is more or less as before, except that we have added an additional excess return in bonds. Again, we could include an intercept, just for the sake of testing whether it is zero as required by the CAPM. This extended model is actually similar to something called the *APT model*, but it's unnecessary to develop the comparison here.

From the interpretive point of view, the coefficient  $\beta_i$  measures the marginal effect on the subject stock #  $i$ , of an increase in the market excess return. The other variable—namely, the bond excess return—is imagined to be held constant. Likewise, the parameter  $\gamma_i$  gives the marginal effect of an increase in the bond excess return. Collectively, the two excess return variables on the right side of the regression equation are called the *independent variables*, while the left hand variable (in this case our stock excess return) is called the *dependent variable*.

The least squares theory we described earlier carries over almost entirely unchanged to the statistical fitting of the multiple regression model (8). For each chosen stock  $i$ , the program will compute the estimated parameters  $\beta_i, \gamma_i$  that minimize the sum of the squared residuals. The program will print out the estimated parameters together with their estimated sample standard deviations.

In place of the correlation coefficient as a measure of the goodness of the fit, is a more general statistic called the *multiple correlation coefficient*, conventionally denoted  $R$  or  $R^2$ . It has a similar interpretation in terms of the percentage of the variance in the dependent variable collectively explained by the independent variables, and lies between zero and one, the latter indicating a perfect fit. There is now an additional measure called the *corrected R square*, which is a somewhat artificial way of allowing for the fact that the more right hand variables you add, the more you automatically increase the closeness of fit—that is you diminish the residual variance. Since this is an artifice of more variables, there should be a way of discounting this sort of “buckshot” approach, and this is what the corrected  $R$  is all about – it penalizes the observed fit of the relationship to compensate for the number of right hand variables added.

Finally, you can test the fitted coefficients as before, using the  $t$  score, or if the number of degrees of freedom is fairly large, using the  $z$  score and the standard Normal table. The degrees of freedom are equal to the number of observations

less the number of right hand parameters. Thus, for two right hand regressors, there are three right hand parameters (the intercept as well as the two slope parameters), and the degrees of freedom are  $N - 3$ .

### Example 7.5

Let's try it. To perform a multiple regression using Excel, the data for the independent variables must sit in contiguous columns. This means it will be necessary to copy the "market returns" data into a new column next to the US Govt bond data in column P. With the cell pointer on [Q1], type: `=A1` into the cell. Then copy cell [Q1] all the way down the Q column as far as cell [Q60]. (Make a few random checks to see that the data is the same as that in column A.) Your data is now set up for multiple regression.

Once again, slide the cell pointer over to cell [C1] and in the menu toolbar, click on *Tools, Data Analysis...*, and then select *Regression* from the "Analysis tools" box, and click on [OK].

If you are using GM, nominate [`$C$1:$C$60`] as the Input for the Y Range, and [`$P$1:$Q$60`] as Input for the X Range. Then select the following:

Labels

For Output Options, select:

Output Range (and specify an area clear of the data, such as `$T$95`)

For Residuals, select:

Residuals                       Line Fit Plots

And click [OK].

**Table 7.2 Multiple regression output for General Motors**

## SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.532882707
R Square	0.28396398
Adjusted R Square	0.258391265
Standard Error	0.083042793
Observations	59

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	2	0.153151167	0.0765756	11.104178	8.67314E-05
Residual	56	0.386181911	0.0068961		
Total	58	0.539333078			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>
Intercept	-0.00488087	0.011125588	-0.4387069	0.6625619	-0.02716809	0.0174063
Bond	0.383079606	0.506508296	0.7563146	0.4526307	-0.63157776	1.397737
Market	1.094071078	0.234284002	4.6698497	1.929E-05	0.624744137	1.563398

**Results**

Table 7.2 represents an edited version of the Excel output for GM as dependent variable against the ten-year treasury bond return and the market return as independent variables, in each case measured as excess returns relative to the bill rate. You will see that once the market return is there, the bond rate adds very little. Its  $t$  value against the zero alternative is statistically insignificant. The multiple  $R^2$  goodness of fit statistic is only marginally better than the earlier  $r^2$  we estimated from the market return alone. Thus, the fixed interest market in the US appears to add little, if anything, to the explanation of variation in common stock returns, at least for GM.

1. Do scatter diagrams for the following (the first variable named should be on the vertical axis):
  - (i) Bond returns versus CD returns.
  - (ii) Bond returns versus stock market returns (the S&P500).
  - (iii) Stock market returns versus CD returns.

Which of the various pairs appears to be the best fit?  
How can you explain different slopes between the diagrams?

2. (a) How might regression techniques be used for purposes of prediction?
  - (b) If you use regression models for this purpose, what are the likely sources of error in your forecast? Assume that you have to first fit the model off an initial sample period.
3. Aptitude test: Can you be a stock analyst?
  - (a) Look at the company names associated with the data in the data file. Bearing in mind what you know about the way that stock returns vary according to the cyclicity of the stock (e.g. whether the stock varies with the general market, and if so, whether “more so”), give the companies a provisional ranking as to where you think their betas would be, relative to one another. (To use the appropriate terminology,  $b > 1$  is said to be an “aggressive” stock, and  $b < 1$  a “defensive” one.)
  - (b) Check your economic intuition by actually computing the time series beta for each of the stocks given in the data file. List the betas in ascending or descending order, and see how closely this corresponds with your *a priori* gut feeling about them.

(Of course, to be a good stock analyst, you also have to be able to constantly change your mind without batting an eye!)
4. As a CAPM exercise, you can take a different tack and look at the beta for the government bond index against the share market. Theory suggests that at least for a pure discount bond, there should be such a beta. You could certainly try to compute a beta for bonds overall, against the stock market, using the data from the web site. Thus, get the excess return by subtracting the CD returns off the bond returns and regress this on the market excess return. Is

the bond beta greater or less than one, and how might this tie in with your prior expectations?

5. The ordered mean difference (OMD) was developed\* as a technique in performance measurement that measures how well a given fund's return does against a benchmark return. Essentially, it captures the insight of the song "Whatever you can do, I can do better," from Rodgers and Hammerstein's *Annie Get Your Gun* (Feel free to hum along.) The following Excel exercise utilizes the menu tools and wizards.
  - (i) Take one of the stocks in the disk. Tabulate the returns on the stock together with the return on the market at the same date. Now reorder the observations according to increasing values of the benchmark. Next, compute the running mean of the differences, which is the progressive moving average as you move down. The table below illustrates the calculation. You end up with a function that you can plot against the benchmark return (OMD as a function of R).
  - (ii) If the OMD function lies wholly above the horizontal axis (all values are positive) this means that your chosen stock (or fund, etc.) is dominant over the benchmark (market, in this case). Sell the market and buy the stock. More commonly, in a group of stocks (such as the ones on your data disk) some will have OMD schedules that slope upwards (aggressive) others that slope downwards (defensive). If the CAPM hypothesis is correct, all should cross the horizontal axis at the same point. Test the CAPM using the stocks from your data disk by plotting all their OMD schedules on the same diagram.

**Table 1 Sample OMD calculation**

r values	R values	Reordered by R		running mean difference	
		r	R	r-R	OMD
0.032	0.011	-0.021	-0.045	0.024	0.024
0.048	0.047	-0.026	-0.030	0.004	0.014
0.028	0.018	0.004	-0.005	0.009	0.012
0.004	-0.005	0.017	0.001	0.016	0.013
0.027	0.008	0.027	0.008	0.019	0.014
-0.026	-0.030	0.032	0.011	0.021	0.016
-0.021	-0.045	0.028	0.018	0.010	0.015
0.017	0.001	0.048	0.047	0.001	0.013

\*Bowden, R.J. (2000), "The ordered mean difference as a portfolio performance measure," *Journal of Empirical Finance*, 7, 195-223.





---

# Eight

## Introduction to Stochastic Processes

Most variables that we deal with in finance have some sort of time dimension. For example, think of stock prices, bond prices, or interest rates. In making any sort of portfolio or risk management decision, the essence of the matter is what the current values are and what they are likely to be in the next period or at the end of the decision horizon. Thus, there are at least two random variables to consider: the value of our stock price now (or whatever the asset), and its value in the future. One could use the same name, for example “S” for the stock price, and two different time suffixes, such as  $t$  and  $t+1$  so that the two random variables are  $S_t$  and  $S_{t+1}$ . If you think about it, the information that you use for your guessing game about  $S_{t+1}$  would certainly include past values such as  $S_t, S_{t-1}, S_{t-2}$  and so on. In other words, we are dealing with a potentially infinite set of random variables—the values of the stock price at all times past, present and future. A *stochastic process* is the name given to such an infinite set of random variables, with time forming the indexing set with representative element  $t$ . Moreover, time can be treated as a continuous element, which turns out to be very convenient.

Stochastic processes form the underlying mathematical framework of modern finance. In this chapter, we aim to give you an idea of the nature of this framework without delving too deeply into the mathematical subtleties, which at times can become very complex. We shall draw your attention to some of the applications – and the jargon—so that you’ll have some idea of what the quants are talking about.

## 8.1 Notation and definitions

ARCH	An econometric model for moving volatility, that says the current volatility is a function of the past volatility plus the absolute value of the equation error or disturbance this period.
Drift	The expected value of the change in a variable, conditional upon past information, appearing as the mean term in an Ito process.
Implied volatility	The volatility ( $\sigma$ ) obtained when the market price of an option is set equal to the price given by a particular option-pricing model.
Ito process	A model in instantaneous time that explains the change of a variable in terms of a mean (or drift) term, possibly changing in time, and a purely random effect (the volatility term) generated by a normally distributed innovation or white noise process.
Martingale	A stochastic process in which the expected value of the future value of the variable, given all past and current information, is simply equal to its current value. On average, the variable stays the same, though its variance can explode.
Random walk	Special case of a martingale in which the random effects have the same distribution through time, so one can write the current value of the variable as equal to the immediate past value plus a noise term of zero expectation.
Stochastic process	A sequence of random variables, usually referring to successive values of the same named variable, and assumed to obey some sort of causal law connecting the values in successive periods.
Stochastic variable	Another name for a random variable, one whose values as of any future point in time are not known now.

## 8.2 Expectations

One is accustomed to see surveys of consumer “expectations” and the like, but the word will be used in a more precise sense here. The expectation of a random variable means its mathematical expectation, which is simply its arithmetic mean. Thus, if  $x_t$  is a random variable representing the value of a magnitude  $x$  at time  $t$ , we write its mean as  $E[x_t]$ , or just  $E x_t$  for short.

We already have a rough idea of the mean from Chapter 6. The arithmetic mean of a random variable is the average value obtained by the sum of the possible values, each weighted by the frequency, or probability, attached to it. In the context of a stochastic process, the mean  $E x_t$  at time  $t$  represents the *unconditional expectation of  $x_t$  at time  $t$* . Imagine that there is a glorious time machine that enables you to rewind and repeat the whole history of stock market prices, over and over again. Each time you do it, there is some randomization device that says that you get a new time path (called a *sample function* in the literature). The average at any specific time  $t$  of all such realized values is the unconditional expected value at time  $t$ .

You might think that all this business of imagining re-runs is academic, because it is an inescapable part of our human condition that we cannot re-run time. And you would be right – the unconditional expectation plays only a very limited role in financial mathematics. Indeed, many of the models used to generate prices and interest rates do not result in a finite unconditional expectation, so that the unconditional expectation does not even exist.

A more informative approach would be to ask the following: Suppose that I know that the stock price of a given stock yesterday was \$10. Would this affect my expectation of where the stock price tomorrow would be? Obviously, the answer is yes. Knowing today’s stock price is important information that conditions my probability assessments of the value of tomorrow’s. Of course, there might be other information available today that also affects my probability assessments of tomorrow’s stock price – I might have learned that the NAPM (National Association of Purchasing Managers) index has turned downwards. I might therefore extend the information set  $I_t$  available to me today to include the state of the NAPM index, as well as yesterday’s price, and clearly, both are useful in predicting tomorrow’s stock price. One could regard the unconditional expectation of the stock price as the mean, with no information of any kind available. But if there is information available, one could assert that in general,

$$E S_{t+1} | I_t \neq E S_{t+1}$$

The vertical bar is the conditioning sign. It says that the mean is conditioned by knowing the information that appears after bar. If the information set comprised  $S_t$  and  $M_t$ , the latter representing the state of the NAPM index, one could write the conditional expectation as  $E S_{t+1} | S_t, M_t$ .

The conditional mean, or conditional expectation, plays the central role in much of finance theory – and also in practice, as the extended example in the next section will show.

### 8.3 Hedging

As you know by now, hedging refers to minimizing the variation in the value of a given portfolio. For instance, if you have a receivable at the end of the coming period that is denominated in Canadian dollars (CAD), you might decide that you don't want the currency exposure. You augment the implied portfolio of CAD with a forward position in CAD, that precisely cancels out the exposure, minimizing – right down to zero – the variation in the value of your portfolio in US dollars. Similarly, if you have an exposure to the price of fine wool, you might want to augment it with wool futures. In the first case, you can reduce the exposure right down to zero. In the second case you can't, because wool futures are based on a composite grade and the price of the futures contract will not track precisely along with the price of fine wool. In the latter case, one would say that there is residual *basis risk*. We can also say that basis risk means that hedging is imperfect.

Suppose that your exposure is to an asset of value  $Y$ , and you have another asset of current value  $X$  available as a hedging instrument. The latter is not a perfect hedge. However, you do know that changes in their values are linearly related. Thus, the following would hold for any period  $t$ .

$$\begin{aligned} \text{Let: } y_t &= Y_t - Y_{t-1} \quad (\text{change in } Y, \text{ the position value}) \\ x_t &= X_t - X_{t-1} \quad (\text{change in } X, \text{ the hedge instrument value}) \end{aligned}$$

$$\text{then: } y_t = \beta x_t + \varepsilon_t \quad (1)$$

We are specifying that  $E y_t | x_t = \beta x_t$ . In other words, once I know  $x_t$ , my expectation of  $y_t$  is equal to  $\beta x_t$ . This means that  $E \varepsilon_t | x_t = 0$ . The term  $\varepsilon_t$  represents a random disturbance that separates  $y_t$  from  $x_t$ , because otherwise, the two would be essentially indistinguishable. One would simply be a rescaled version of the other. The disturbance has the property that knowing  $x_t$  provides no information about  $\varepsilon_t$ , so that the latter can be regarded as irreducible variation in  $y_t$  relative to  $x_t$ .

In the context of hedging, imagine that we are at time  $t$  looking ahead to time  $t + 1$ . Our exposure is to the change  $y_{t+1}$  in the position Y. We have a hedging instrument X available whose change in value  $x_{t+1}$  at time  $t + 1$  is also unknown at time  $t$ . But we can nonetheless use it to hedge the unknown change  $y_{t+1}$ , taking advantage of the relationship (1).

To set this up, to my existing exposure at time  $t$  of 1 unit long in Y let me go  $q$  units short in X, which means a negative position. The change in value in my hedged position between times  $t$  and  $t + 1$  is

$$v_{t+1} = y_{t+1} - q x_{t+1}$$

We need a loss function for the discrepancy. We shall use the *mean square error loss function*, which is defined here as  $E v_{t+1}^2$ ; the expectation is once again conditional on information available to us at time  $t$ , so that technically we really ought to write it as  $E v_{t+1}^2 | I_t$ . Note that while the loss function penalizes large deviations, as we would expect, it is two sided, penalizing value gains as well as value losses. The latter property is characteristic of pure hedging strategies. For example, if you hedge with futures, you can protect yourself against losses but you also dip out on potential gains. Pure hedging strategies may therefore be defined in terms of position sterilization.

Now if we insert (1) into this loss function, we get the minimand (that is, the value we wish to minimize):

$$E v_{t+1}^2 | I_t = E [(\beta - q) x_{t+1} + \varepsilon_{t+1}]^2 | I_t$$

The problem is to choose  $q$  to minimize this expression. It is easy to show that the minimizing choice is just:

$$q = \beta$$

So the optimal hedge for every unit of exposure in Y is to go short  $\beta$  units of X. When you do this, the irreducible residual MSE is just  $\sigma_\varepsilon^2 = E \varepsilon_{t+1}^2$ . This is the *basis risk*; we will see how to interpret it in the following example.

Now, how about the empirical implementation? Astute readers will have noted that Equation (1) looks very much like a regression equation, and so it is. We have in effect shown that the best hedge, with the MSE loss function, is derived in terms of the regression equation of the variable or position to be hedged upon the hedge instrument. In fact, the optimum hedge is always the conditional expectation of the thing to be hedged, given the hedge instrument, and OLS regressions are a means

of identifying the conditional expectation. Then you just use the fitted coefficients to derive the hedge portfolio. Finally, the regression estimate of the residual variance gives us the minimum residual MSE, in the sense that its square root will give us one standard deviation of the likely monetary loss remaining after the hedge. Note that as a regression property,  $\sigma_\epsilon^2 < \sigma_y^2$ , which means that the hedged variance is always less than the variance of the unhedged position.

These principles can apply to other situations. For instance, you could think of improving the *goodness of fit* by having a hedge comprised of two instruments rather than just one. To guide you to the right portfolio proportions, you would do a multiple regression of the (changes in the) position to be hedged on the changes in the values of the two hedge instruments, as the independent variables in the regression.

Here's an example of how it works:

### Example 8.1

#### Metal and futures hedge

Let us suppose that you are working in the treasury of a major munitions producer that uses input of copper and lead for the manufacture of bullets, shotgun cartridges and other purposes.

You have daily metals price data for the past few years. (You can find a subsection of this data, as a file entitled "coprlead.xls," on the Authors Academic Press website at [www.AuthorsAP.com](http://www.AuthorsAP.com)). The data includes daily prices of the following:

1. Spot price of copper expressed in USD per ton
2. Futures price of copper (three-month contract) expressed in USD per ton
3. Spot price of lead expressed in USD per ton

If you take a look at the data, you will see that sometimes the futures price is above the physical price and sometimes the reverse is true. The differential between the futures and the physical price is governed by several things, including interest rates, cost of storage, commodity shortages, etc.<sup>1</sup>

<sup>1</sup> For a well-known case where these factors changed dramatically, see one of the many accounts of Metallgesellschaft, such as Edwards, Franklin R, and Michael S. Canter "The collapse of Metallgesellschaft: unhedgeable risks, poor hedging strategy, or just bad luck?" *Journal of Applied Corporate Finance*, 8, 86-105.

Your main exposure to price risk results from your need to purchase copper, but you have some direct exposure to lead prices as well. While it is possible to obtain both lead and copper futures for hedging purposes, the relatively small amount of lead that you buy makes the use of lead futures problematic. It is a nuisance to have to undertake the management, systems development, input, and settlement of two different types of futures contracts in addition to the various dates for each contract. In view of this problem you have decided to explore the possibility of using copper futures for the hedging of both types of metal exposure.

**Table 8.1 Regression output:  $\Delta$ Copper on  $\Delta$ Copper futures**

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.920276453
R Square	0.84690875
Adjusted R Square	0.846755198
Standard Error	17.37170699
Observations	999

ANOVA					
	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	1664433.317	1664433	5515.456	0
Residual	997	300870.8751	301.7762		
Total	998	1965304.192			

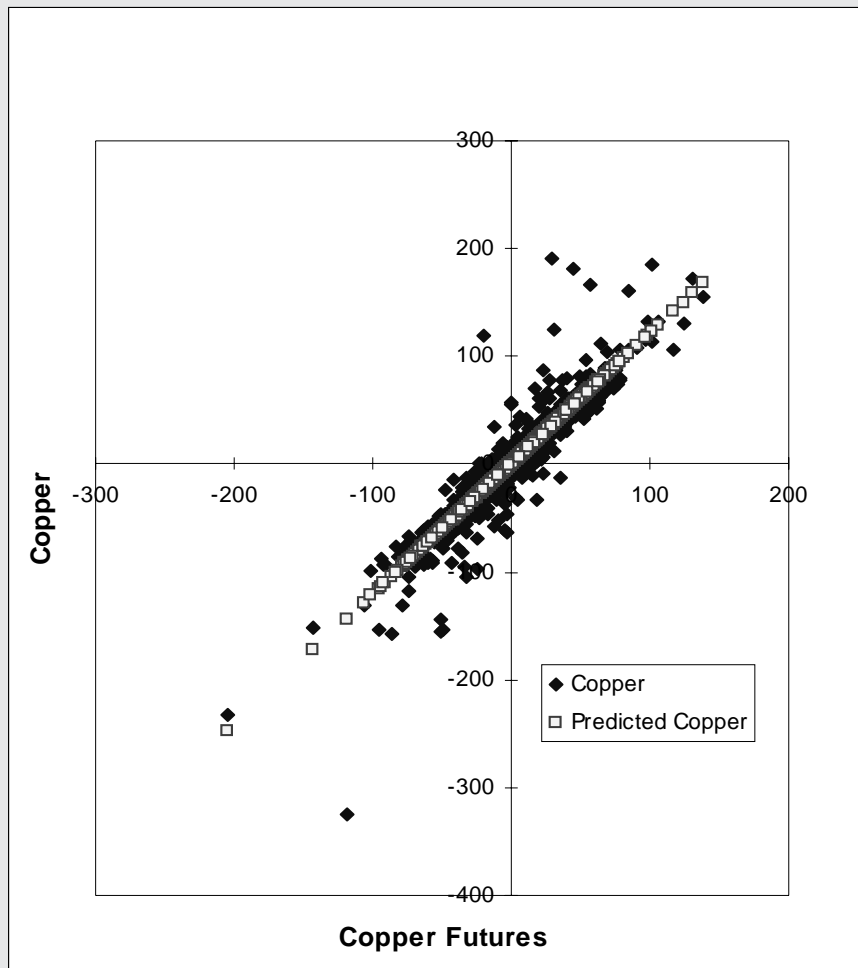
  

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>
Intercept	0.097014977	0.549875199	0.176431	0.859991	-0.9820293	1.176059
Copper Futures	1.210385307	0.016297948	74.26611	0	1.178403127	1.242367

RESIDUAL OUTPUT

<i>Observation</i>	<i>Predicted Copper</i>	<i>Residuals</i>
1	-39.2211414	-23.7758586
2	-38.6813095	-12.5166905
3	86.1820388	-6.0820388
4	-37.5943835	1.7833835
5	-34.7766065	-8.4223935

For the first part of the analysis, you will need to find the number of copper futures contracts required to hedge your copper exposure on a per ton basis (assuming, for simplicity, that you hedge your requirements no more than 90 days in advance of the physical purchase). Note that each contract is for five metric tons of copper. To solve this problem, you will start by running a regression in Excel with the *daily change in the price of copper futures* as your X input and the *day to day change in the price of spot copper* as the Y input. This has been run already with 999 observations, resulting in the output displayed Table 8.1 and Figure 8.1.



**Figure 8.1**  $\Delta$ Copper spot regressed on  $\Delta$ Copper futures

As you can see from Table 8.1, there is clearly a strong linear relationship (as you would expect) between changes in the price of copper and that for copper futures, as indicated by an R Square of 0.847. You can also see how closely the observed points lie to the fitted line in Figure 8.1 below. The regression coefficient of 1.21 tells us that for each change per ton of (USD) \$1 in the price of copper futures, the spot price is expected to change by (USD) \$1.21. So for every ton of copper we wish to purchase in the coming quarter, we need to buy 1.21 tons equivalent in copper futures contracts. As a copper futures contract is for five metric tons, our requirement for hedging a (metric) ton of copper will be:



$$\text{number of copper futures contracts} = 1.135 / (5 \text{ tons}) = 0.227$$

(a bit less than 1/4 contract)

The Standard Error of 17.37 indicates that with 95% probability, we expect to lose (or gain) no more than \$34.74 per ton on our overnight hedging strategy (2 x \$17.37).

For the second part of the exercise, run another regression using the one-day change in the price of lead as the X input and the one-day change in the price of copper futures as the Y input. Again, how many copper futures will you need?

The results of part two, using 999 observations, are given below as Table 8.2 showing the summary regression output, and Figure 8.2 showing a scatter graph with fitted regression line.

**Table 8.2 Regression output:  $\Delta$ Lead spot on  $\Delta$ Copper futures**

SUMMARY OUTPUT

<i>Regression Statistics</i>	
Multiple R	0.884253491
R Square	0.781904237
Adjusted R Square	0.781685485
Standard Error	42.54337127
Observations	999

ANOVA

	<i>df</i>	<i>SS</i>	<i>MS</i>	<i>F</i>	<i>Significance F</i>
Regression	1	6469419.282	6469419	3574.386	0
Residual	997	1804508.624	1809.938		
Total	998	8273927.906			

	<i>Coefficients</i>	<i>Standard Error</i>	<i>t Stat</i>	<i>P-value</i>	<i>Lower 95%</i>	<i>Upper 95%</i>
Intercept	0.465789518	1.34664629	0.345889	0.729499	-2.17679369	3.108373
Copper Futures	2.386289291	0.039913733	59.78617	0	2.30796482	2.464614

RESIDUAL OUTPUT

<i>Observation</i>	<i>Predicted Lead</i>	<i>Residuals</i>
1	-77.050432	-49.710768
2	-75.986147	-28.203553
3	170.183457	-83.491557
4	-73.843259	18.085459
5	-68.287978	-48.714822
6	-26.511211	47.436711
7	-77.358263	12.239363

An R Square this time of 0.78 again indicates a close relationship in the overnight price changes, this time between lead and copper futures. Given a standard error of 42.54, our hedge will not be as safe as it was for copper, but we can still use these contracts with 95% certainty that we should not lose (or gain) on an overnight hedge more than \$85.08 per ton of lead hedged with copper futures ( $2 \times \$42.54$ ). The regression coefficient of 2.386 indicates that the per ton price of lead moves by \$2.39 (in USD) for each one dollar (USD) move in the price of copper futures. Hence, to hedge our next quarter's purchase of lead we have:

copper futures requirement =  $2.386 \times \frac{1}{5} = .4772$  (about half a contract)

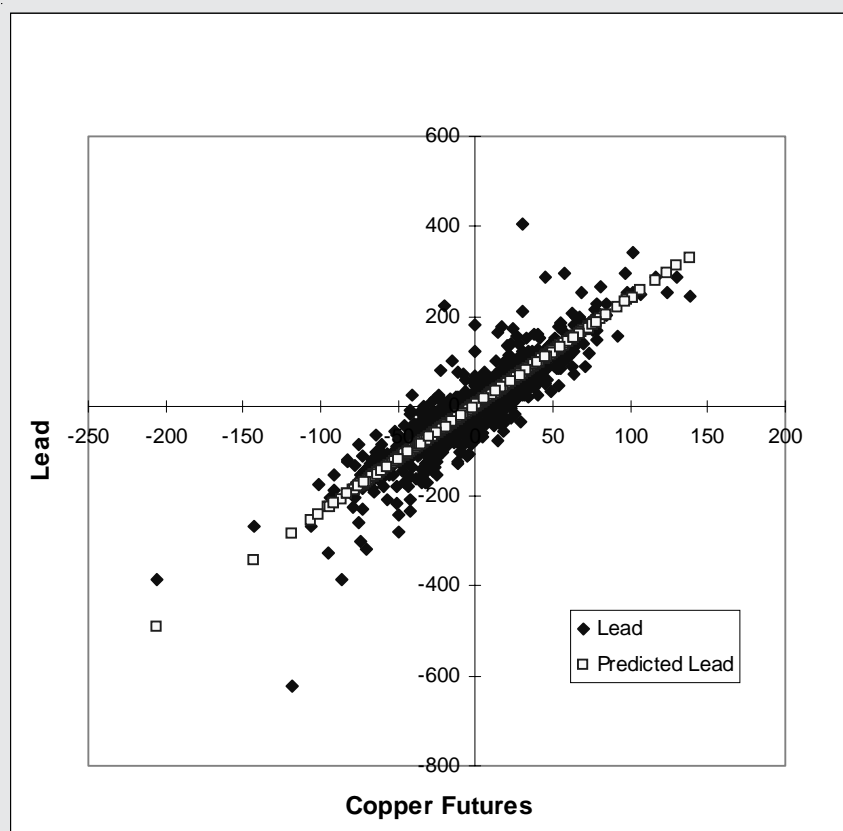


Figure 8.2  $\Delta$  Lead spot regressed on  $\Delta$  Copper futures

## 8.4 Random walks and Ito processes.

Now we shall look at some of the modeling processes for financial variables. Understanding something about the structure of these processes will help you to comprehend the assumptions underlying derivatives pricing and we'll say a little more on that later.

From time to time, you may see claims that stock market prices follow random walks. The claim itself is not quite correct, even under the kind of assumptions that the authors have in mind, but there is no doubt that random walk elements are involved. Indeed, the ultimate building blocks of most financial processes are made out of something very closely related, so random walks are a good place to start.


Consider the process  $(B_t)$  which obeys the following recursion:

$$B_t = B_{t-1} + \varepsilon_t \quad (2)$$

The value of  $B$  at time  $t$  is obtained by adding the value at time  $t - 1$  to a random variable  $\varepsilon_t$ , for which all successive values are uncorrelated; in other words,  $\varepsilon_t$ ,  $\varepsilon_{t-1}$ ,  $\varepsilon_{t-2}$ , ... are all independent of each other. These are often described as *innovations*. Think of them as “news” variables—news is only news if each element is unexpected, or could not be predicted from the previous elements. Sometimes the  $\varepsilon_t$  is called a *white noise* series, because the series  $\varepsilon_t$  has no recognizable structure, like static on a radio (as distinct from music...though sometimes it's hard to tell the difference!)

As its name suggests, a random walk as defined in Equation (2) wanders aimlessly. To get this period's value, you simply take last period's value and add something completely random. The result is that if you start at time  $t = 0$ , and keep doing this, after some time, the possible values of  $B$  form a pattern that is ever more diffuse; that is, the variance of  $B$  tends to infinity. After a long time, you have no clue where the value of  $B$  is likely to fall. This is technically called the *unit root property* and is an example of a nonstationary series. Figure 8.3 illustrates.

This is an example of a process where there is no stationary distribution. You can see this intuitively. Suppose you started continuously generating values of  $B$  and putting them in histogram-type relative frequency boxes. Because the process continues to spread out – its variance expands forever – you would have to keep adding histogram boxes at each end, collecting proportionately fewer and fewer in any given interval.



**Figure 8.3 Random walk sample paths**

A generalization of a random walk is a *martingale*. This is a process (using  $B$  again) such that if  $I_t$  is an available information set at time  $t$ , then  $EB_{t+1} | I_t = B_t$ . You can see that a random walk is certainly a martingale, as the information set available at time  $t$  for the random walk would be imagined to include the value of  $B$  at time  $t$ . Thus, to predict a martingale, the best you could do is use the current value as a predictor. The rest is just “news” that you cannot predict.

As mentioned above, it is often supposed that under efficient markets, a stock price follows a random walk. The idea is that the existing stock price impounds all available information, so is itself the best predictor of next period’s stock price. This is half right. The correct statement is that if you plough back all dividends and other elements to form an accumulation fund (see an investment text for a definition), then the successive values of this fund, discounted back to some starting point, form a martingale (provided investors are risk neutral).

Now we go on a step further. Suppose we had a stock whose price  $S_t$  obeys the following:

$$S_{t+1} - S_t = \mu S_t + \sigma S_t \varepsilon_{t+1} \quad (3)$$

Let’s do several things with this equation. First, we would normalize the variance of the (normally distributed) innovations terms  $\varepsilon$  to be unity. Then:

- (i) The left hand side is the change in the price of the stock.
- (ii) The first term of the right hand side is known at time  $t$ , which we shall take here as the current time. It says that there is a predictable component to the change in the stock price, which corresponds to a percentage growth  $\mu$ . This term is called the *drift* term of the process.
- (iii) The second term on the right hand side contains the news to transpire between  $t$  and  $t + 1$ . We conventionally normalize the variance of the innovation itself (that of  $\varepsilon$ ) to be unity. The term  $\sigma$  then becomes the instantaneous standard deviation of the random component of the proportional change in  $S$ . It is also called the *volatility*.

Equation (3) could be rewritten as an equation in the proportional change:

$$\frac{\Delta S_t}{S_t} = \mu + \sigma \varepsilon_{t+1} \quad , \quad \text{where} \quad \frac{\Delta S_t}{S_t} = \frac{S_{t+1} - S_t}{S_t} .$$

If you divide both sides of (3) through by  $S_t$  as shown, then you can clearly see that the drift parameter  $\mu$  and the standard deviation parameter  $\sigma$  refer to the proportional changes. Now let's take yet another step. We can think of a process  $B_t$  whose increments are equal to the given innovations process,  $\varepsilon_t$ . That is,  $\Delta B_t = B_{t+1} - B_t = \varepsilon_{t+1}$  is a sort of cumulated state of the world, the sum of all past news.

Then we can rewrite (3) as:

$$\Delta S_t = \mu S_t + \sigma S_t \Delta B_t \quad (4)$$

Finally, convert to continuous time by imagining that the time intervals for our discrete time  $t$  process get smaller and smaller, and infinitesimal. As Appendix A5 on calculus shows, we imagine that the finite increments  $\Delta S_t$  are replaced by the infinitesimal increments  $dS_t$ . In this case, (4) is written schematically in the form:

$$dS_t = \mu S_t + \sigma S_t dB_t \quad (5)$$

The term  $\mu$  is called the instantaneous *drift* of the process  $S$ , while the term  $\sigma$  is called the instantaneous *volatility* of the process. The process  $S_t$  is an example of an *Ito process*.

Equation (5) can also be read as follows:

If we imagine the division of time into infinitesimally small intervals  $t, t + dt$ , then the conditional expectation of  $S$  at time  $t + dt$ , given information  $I_t$  available at time  $t$ , is given schematically by:

$$E ( S_{t+dt} - S_t ) | I_t = \mu S_t \text{ times the interval length } dt$$

In other words:

$$E \frac{(d S_t)}{S_t} = \mu dt .$$

So  $\mu$  is the instantaneous expected rate of increase in the stock price.

Along with an instantaneous mean, the stock price  $S$  at time  $t + dt$  will have an instantaneous variance. In this case, things are constructed so that the variance of the proportional change in  $S$  is equal to  $\sigma^2$  multiplied by the interval length  $dt$ .

You can use this property to estimate  $\sigma^2$ . Suppose your unit of time is one day. Then the  $\sigma^2$  is simply the variance of daily proportional changes in the physical. Alternatively, if you assume the unit of time is one year, then  $\sigma^2$  is the annual variance. Because these are variances of proportional changes, the corresponding standard deviation  $\sigma$  comes out with the dimension of percentage. So if the time period is annual, the  $\sigma$  is the standard deviation of annualized proportional changes. See Example 8.2 below.

Finally, both  $\mu$  and  $\sigma$  can depend on time, and even on the value  $S_t$ . For example, we could specify:

$$\mu_t = k ( S^* - S_t ) ,$$

where  $S^*$  is some constant value and  $k$  is a constant representing a speed of adjustment. This sort of formulation is called an *error adjustment mechanism*. It says that there is some target level of the stock price  $S^*$ , and the stock price can be expected to drift upwards if the current value is less than the target, or downward if the current value is greater than the target. In addition, the second term on the right hand side is the unforeseeable random volatility component, which is essentially unpredictable.

### Example 8.2

#### Example of an Ito process

Suppose a stock that pays no dividend has a volatility of 29% per annum and an expected return of 12% per annum (by convention, we assume continuous compounding). We shall operate with the year as the unit of time. So we can set  $\mu = 0.12$  and  $\sigma = 0.29$ . The Ito process for the stock price is:

$$dS/S = 0.12dt + 0.29dB$$

If  $S$  is the price of the stock at some point in time and  $\Delta S$  is the increase in the stock price over some small interval in time,  $\Delta B = \varepsilon \sqrt{\Delta t}$ , and

$$\Delta S / S = 0.12 \Delta t + 0.29\varepsilon \sqrt{\Delta t}$$

where  $\varepsilon$  is a random drawing from a standardized normal distribution. Consider  $\Delta t$  as a week ( $7/365 = .019178$ ) and suppose that the beginning stock price is \$75.00. In summary,  $S = 75$  and  $\Delta t = .019178$ . Thus the change in  $S$  is:

$$\Delta S = 75(0.12 \times 0.019178 + .29 \times \varepsilon \times \sqrt{.019178})$$

or 
$$\Delta S = 75 \times (0.00230 + \varepsilon \times .04016)$$

$$\Delta S = 0.1726 + \varepsilon \times 3.012$$

showing that the price increase can be regarded as a random drawing from a normal distribution with mean of \$0.1726 (17.26 cents) and standard deviation of \$3.01.

## 8.5 How Ito processes are used

Ito processes of the kind we have just described are a vital tool for theoreticians in deciding how to price derivatives correctly. In this context, to price things correctly means that nobody can arbitrage against you by finding that you have systematically under priced or overpriced the derivative you are offering or trading; therefore, we refer to the process as finding a *no-arbitrage* price. The details are often highly mathematical, so if your math is a bit incomplete, you may have to skip this subsection.

The typical sort of derivative is something like an option or a forward payoff maturing at some definite time  $t$  in the future. So at some future time  $T$ , it will have a payoff that is a specified function—say,  $g(S_T)$ —of the future physical price  $S_T$  at that time. By the “physical” we mean what you are tying the derivative to – for example the physical for stock options is the stock price, for bond options, it’s the bond price, and so on. In the case of a call option with strike price  $X$ , for example, the function  $g$  is defined by:

$$\begin{aligned} g(S_T) &= \begin{cases} S_T - X & \text{if } S_T > X \\ 0, & \text{otherwise.} \end{cases} & (6) \\ &= \max(S_T - X, 0) \end{aligned}$$

The payoff is a promise to pay a formula amount that is contingent upon whatever the final physical price turns out to be. At time  $t < T$ , such a promise will have a market value, and the trick is to find out just what it ought to be.

At first sight, this might appear to be like deciding how long is a bit of string – the promise could be worth anything, depending upon what people are willing to pay for it, who is in the market, and so forth. However, there *is* a unique correct price. The reason is that the derivative provides no new information. It is purely a piggyback instrument. Since it provides no new information, you ought to be able to replicate all its price movements in the interim through to time  $T$ , by means of a position in the underlying physical, plus a risk-free asset such as a call account (often bonds are considered an example of this, but we know by now that bonds are hardly risk-free!). It is this principle that enables us to derive the equation that the derivative price should satisfy, and, by solving that equation, the theoretical formula for the price of the derivative as well.

Suppose that the current price of the option is some unknown function  $C$  of the current stock price  $S$  and time  $t$ .

$$Y_t = C(S_t, t), \quad t < T \quad (7)$$

According to a result called *Ito's lemma*, you can say that  $Y_t$  must itself obey an Ito process, and you can specify what the mean (i.e. the drift term) of this process must be in terms of the partial derivatives of the function  $C$ . (See Appendix A5 for a discussion of partial derivatives.)

Then you construct a replicating portfolio in terms of some – as yet to be determined – portfolio amounts  $a_t, b_t$ , so that:

$$Y_t = a_t S_t + b_t A_t, \quad (8)$$

where  $A_t$  is the price of the call account, a notional construct in which you assume that \$1 has been invested at  $t = 0$ , and allowed to accumulate at the risk free rate until the chosen time  $t$ . In other words, with continuous time accumulation as in chapter 2.4,  $A$  is equal to  $e^{rt}$ , where  $r$  is the risk-free rate.

By matching up the amounts  $a_t, b_t$  in (8) with the drift term in the Ito process for  $Y_t$  as derived from (7), you can write down a partial differential equation that the unknown function  $C$  must satisfy. Then you add the terminal or boundary condition (6), and—*presto!*—out comes the theoretical price of the option. It turns out that this price depends on a number of factors: the current stock price  $S_t$ , the volatility  $\sigma$  of the  $S$  process, the interest rate  $r$ , the strike price  $X$ , and the time  $t$ , via the time to maturity.



**The Black-Scholes formula**

The Black-Scholes model derives the following equation for the price  $c$  of a call option:

$$c = SN(d_1) - Xe^{-rT} N(d_2) \quad (9)$$

where:  $N(d)$  is the value of the Normal distribution function at the point  $d$ ;  $d = d_1$  or  $d_2$ . The other symbols are defined as follows:

$S$  = current price of the underlying physical

$X$  = strike price of the option

$T$  = time to maturity

$r$  = current risk-free interest rate

$$d_1 = \frac{\ln(S/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} ; \quad (9a)$$

$$d_2 = \frac{\ln(S/X) + (r - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T} ; \quad (9b)$$

$\sigma$  = volatility of the physical

You will notice in this formula that the price does not depend upon the mean  $\mu$ . In everyday language, investor expectations are irrelevant to the price of the option. This may seem like a startling statement. However, it is simply a reflection of the fact that investor expectations do play a role, but only in the price of the underlying physical. The option itself is purely a play on the physical, and the effects of changing expectations can be hedged away by means of the corresponding effect on the physical (the latter being the hedge).

However, this hedging process does not apply to the volatility  $\sigma$ , which appears unavoidably in the price of the option. You can see intuitively why this must be true. Compare two call options, one with a high volatility and the other with low. Suppose that the option is a call option, which is currently a bit out of the money. In other words, the current physical price is slightly less than the strike price of the option, so you would get nothing if you decided to cash in your chip at this point in

time. Of course you wouldn't decide to cash in now anyway, because though your chip is out of the money right now, it could well be *in* the money at some time in the future. It is definitely worth something because of that possibility; hence, the option price is by no means zero. The option with a higher volatility has more chance of getting back in the money next period than does the one with the lower volatility, so it makes sense that it is worth more.

Nevertheless, you do have to be able to form an estimate of the physical's volatility, and we shall discuss this below. In the meantime, we note that if we knew what the option was currently trading for, then we could use an equation such as (9) to reverse the process and solve for the implied volatility in terms of the current price of the option. This process is called *backing out the implied volatility*.

These days, theoreticians use a somewhat simpler approach, which is based on converting the underlying process  $S_t$  into a different world where everybody is risk neutral. In such a world, all assets would have exactly the same expected return; this would be the same as the risk-free rate  $r$ . Here, you can price the derivative as simply the mathematical expectation of the discounted value – at rate  $r$  – of its terminal payoff. To compute that expectation, you use the transformed Ito process—that is, the version that would hold in a risk free world. There are rules that enable you to obtain the modified Ito process, which is called an *equivalent martingale process*.

In practical terms, even the equations for the derivative price are usually rather difficult to solve. The Black Scholes was an exception, though even this eluded researchers for years. As a general procedure for other kinds of derivatives, or other specifications for the driving physical price processes, solving for the options price is often done numerically. This may be achieved by means of a number of approaches. One is to approximate the Ito process in terms of a simpler *Binomial process*. In this approach, the possible variation in  $S$  over a very short time period from  $t$  to  $t + \Delta t$  is constrained to just two possibilities: *up* a certain fixed magnitude, *or down* a certain fixed magnitude. It turns out that you can approximate an Ito process quite well by this method. Arriving at a solution can then be done by means of discrete programming techniques.

From a trading point of view, everything changes once you have a technique for the theoretically correct price. The only unknown parameter in the option price is now the volatility,  $\sigma$ . So deciding what it ought to be is essentially a guessing game. First, one can get a handle on it from the historical record of the variance in the physical prices—we'll look at this in the next section. But on top of that, there is always margin for decision as to what the correct volatility for the coming period ought to be. So trading takes place on the basis of just such a guessing game; in

finance terminology, we say the options market tends to trade on volatility. Players who write options have to think not only about hedging the dependence of the option price as the underlying physical price changes (*delta hedging*), but also hedging as the volatility is judged to be changing (*vega* or *kappa hedging*). Those who try to make money out of options as an uncovered position will do it on the basis of what they think are superior insights as to the  $s$  in the coming period – (these are sometimes called *vega plays* and the practitioners *vegamites*). The term for the activities we’ve been describing is *volatility trading*.

### Example 8.3

#### Backing out implied volatility using Goal Seek

In this example we shall first calculate the Black Scholes option price for a call option on the S&P500 Index. Then we shall assume a different thought experiment in which the option price is set at a certain magnitude and we find out what the implied volatility must be. To do this, we shall use a very useful Excel capability called Goal Seek, which operates by adjusting one of the cells to achieve the result that you want from a given formula.

Here is your base data, and cells we suggest you enter them into:

Data Item	Value	Cell address
Current SPI value	1234	C3
Annual Volatility	0.20	C4
60-day CD rate (annualized)	0.055	C5
S&P500 Divided yield	0.018	C6
Desired Maturity for option (in days)	60	C8
Strike price of option	1250	C9

(So this call option is just a bit out of the money, with no immediate exercise value.) You can enter suitable names for the above data in column A, such as “Current SPI” in [A3], and so on. The unit of time will be annual. The maturity in annual terms will be  $T = C8/365 = 60/365$ . You can take the instantaneous risk free rate  $r$  to be simply  $[C5] = 0.055$ .

#### Task 1: Calculating the Black Scholes price

Set aside cells [C11] and [C12] for  $d1$  and  $d2$  respectively. You will use Formulas (9a) and (9b) to compute them. For the spot or current price  $S$ , we do not use the current SPI value as it appears in [C3], because the

Black Scholes call option price changes a bit when the stock pays a dividend. You don't get the dividend until you have actually called the stock, so in the meantime the lack of dividend entitlement means that you have to discount the current stock price (here the SPI), which does carry a dividend entitlement, by the dividend yield. The simplest way of handling this is to calculate the value of  $S$  as an intermediate value in, say, cell [D10] as:

$$= C3*EXP(-C6*C8/365)$$

and then refer to [D10] when you enter Formula (9a). Note that we are using continuous time discounting here.

When you enter the two formulas for  $d1$  and  $d2$ , it is always useful to enter the cell addresses for the input parameters rather than the values themselves. You will find the formula gets rather long for  $d1$  in particular and you would have to concentrate pretty hard! To make the whole process easier for yourself, break the formula up into intermediate calculations, placing them in cells to the right of the main calculation, as we did above for the value of  $S$ . Watch that left hand and right hand brackets match in number, and are entered in the right places.

When you have finished entering, the results should be  $d1 = -0.04332$  in cell [C11], and  $d2 = -0.12441$  in cell [C12].

Next, enter the call price as calculated from Formula (9) in cell [C14]. You will have to use *normsdist()* function from the Function Wizard (look in *Statistical*), and in the brackets put the cell references [C11] for  $d1$ , and [C12] for  $d2$ . Your final result should come out to \$35.87 as the call option value.

### **Task 2: Backing out an implied volatility**

Now suppose a dealer quoted us an option price of just \$35.00. What could you say about the value of the volatility the dealer was using, assuming that she is also using the Black Scholes formula, but is *not* using your patented volatility estimation technique?

Cell [C14] is the crux of your operation, which brings everything together. Now suppose we put the result of this cell as 35.00, which means that the entry in cell [C4] must be changed. You can try fiddling around on this, but it will take you quite some time to come up with exactly \$35.00 as the outcome.

Goal Seek does the job automatically. Go into *Tools/Goal Seek*, and three options which require entries will pop up. You wish to *set cell* [C14] *to value* 35.00 *by changing cell* [C4]. Once you've done this, the result is produced: an annual volatility of 0.1956, or 19.56% appearing in cell [C4]. This makes sense—because the dealer is quoting a lower price, all else being the same, she must be using a lower volatility.

Don't delete your spreadsheet just yet. You can use it to do problem 5 at the end of the chapter.

## 8.6 Volatility models

The previous section referred to the importance of volatility estimation. The natural starting point is the historical record, even if the trader will subsequently adjust his estimate upwards and downwards by assessing the conditions likely to arise over the coming pricing interval.

A naive estimate of historical volatility is simply the variance of the historical series. A common device is to recognize that volatility can change, but hopefully only slowly, so you adopt a moving sample. Thus, if the series record is monthly over the past 20 years, you could adopt a three-year frame, which would give you 36 observations, and move this along by one month at a time, so that the most recent estimate would be based only on the last 36 months. There are lots of variations on this sort of approach.

A better approach is to combine modeling the conditional mean of the physical  $S$  and the conditional variance. In doing so, we recognize explicitly that the volatility that we need is what will transpire over the coming period, and this is always conditional on the latest information that we have.

To follow this approach, we specify a regression equation. Schematically:<sup>2</sup>

$$S_t = \mu (\text{past } S, \text{ other available variables ; } \beta ) + \varepsilon_t \quad (10a)$$

plus a volatility equation: if  $\text{Var}(\varepsilon_t) = \sigma_t^2$ ,

$$\sigma_t = f(\text{past } \sigma, \varepsilon_t) \quad (10b)$$

<sup>2</sup> Note that in what follows the innovation  $\varepsilon_t$  now has variance  $\sigma_t^2$ , rather than unity.

Equation (10a) is called the mean equation, and in this context  $\mu$  is the conditional mean of  $S_t$  (rather than  $\Delta S_t$ ). The symbol  $\mu(\dots)$  is just a way of saying that the conditional mean  $\mu$  is thought to depend upon a string of variables that are listed inside the brackets, but that we have yet to specify the exact functional nature of the dependence. It says that the physical price is explained in terms of the conditional mean  $\mu$  (which is some function specified by the researcher of past information and any other available variables.) For instance, if you have monthly data, the mean could depend upon dummy or categorical variables to represent the particular month effect. For the definition of seasonal dummies, see an econometrics text.

Equation (10b) is the variance equation, presented with the same generalized notation explained above in connection with the mean. It specifies how the volatility is to evolve. You will notice something about this equation. It is not specified in terms of things that can be directly observed. To get the  $\sigma$  and  $\varepsilon$ 's to put into (10b), you have to estimate them from (10a). More precisely, (10a) and (10b) must be estimated jointly, and this can certainly be done.

Different approaches specify the functions  $\mu(\cdot)$  and  $f(\cdot)$  in different ways. A popular approach is based on the ARCH family, or Auto Regressive Conditional Heteroscedascity. (You can see why the acronym is used!) Since its original development by Robert Engle in 1982<sup>3</sup>, various sons and daughters have been developed (GARCH, MGARCH, IGARCH – AAAARGH!), none of which we will bother explaining here. A typical ARCH specification in finance would be:

$$S_t = \beta_0 + \beta_1 S_{t-1} + \beta_2 x_t + \beta_3 \sigma_t + \varepsilon_t \quad (11a)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \quad (11b)$$

The first equation is the *mean equation*, since the first four right hand variables are intended to capture the conditional mean  $\mu_t$ . The variables  $x_t$  represent any exogenous variables you think might be important. The variable  $\sigma_t$  is the current volatility, which suggests that it might affect the mean (because it feeds back into the pricing of assets).

The second equation is the volatility equation. It says that the residual from the mean equation has a variance that depends upon the latest squared residual available up to time  $t$ . The idea is that if the system gets shocked off the mean path in the last period (a big forecasting error  $\varepsilon_{t-1}$ ), then system volatility should go up. In acronym, version (11) is called a MGARCH (mean generalized ....)

<sup>3</sup> Engle, R.F. (1982) "Autoregressive conditional heteroscedascity with estimates of the variance of United Kingdom Inflation." *Econometrica*, 50, 987-1007.

Equations 11(a) and 11(b) can be estimated by a technique called *maximum likelihood* (ML), which is often used in statistics and econometrics. Essentially it estimates all the parameters (here the betas and alphas) by choosing them to maximize the probability of getting the data series  $S_1, S_2, \dots, S_N$  actually observed. The ML technique is necessary here because the volatility terms  $\sigma_t$  in Equations 11(a) and 11(b) are not observable, and the technique in effect imputes these at each round of the numerical likelihood maximization, which is based on iterating values of the alpha and beta parameters at each step.

The ML technique can be quite tricky to implement, since it relies upon numerical algorithms that may not converge. Also, sometimes the ARCH model simply does not fit. There are alternative specifications to the ARCH family. One of these is due to Schwert<sup>4</sup>, which estimates the current volatility based on the absolute value of the residuals to the fitted mean equation, and uses more or less straightforward regression techniques to execute the necessary steps. For an application to foreign exchange, see Bowden and O'Donovan (1997)<sup>5</sup>, which also describes the technique.

It is important to note that in real life trading, ARCH, Schwert, and other volatility estimation techniques are only the first step. The objective is to begin by estimating the current volatility on this somewhat mechanistic level, and then to look at the result and adjust it according to your own perceptions of the state of the world. For instance, if a Vagon spaceship suddenly appears at the Athens Olympics opening ceremony, it might be worth revising your computed  $\sigma$  upwards (on the off-chance that it is not a publicity stunt).

At any rate, you arrive at a conclusion as to what the volatility should be. Then you look at how the options for the preferred month – or all the months, for that matter – are priced in the market. You can then *back out* the implied options price and compare this with your own view of what the volatility is likely to be. Maybe you have private knowledge from your Sub-Etha Sens-O-Matic of the likely arrival of the aforementioned spaceship. If you think that the market is underestimating the volatility, you can try a bit of volatility trading. Of course, options have many uses, not least of which is hedging portfolios, where they represent a convenient way of protecting the portfolio manager against downside risk in the fund's portfolio of shares or bonds.

<sup>4</sup> Schwert, G.W. (1989) "Why does stock market volatility change over time?" *Journal of Finance*, 44, 1115-53.

<sup>5</sup> Bowden, R.J. and B. O'Donovan (1996). "Financial markets: volatility and policy." In B. Silverstone, A. Bollard & R. Lattimore (eds) *A Study of Economic Reform: The Case of New Zealand*, pp 279-314. Amsterdam: North Holland.

## 8.7 Other time series buzzwords

In many cases, the physical price series  $S_t$  is likely to be non-stationary, which means that its variance and other characteristics are likely to change over time. Actually we have to be a little careful here. By the variance, we mean the *unconditional* variance, which is what you might compute if you took the sample variance of a very large number of observations, ignoring the temporal structure, and treated them as some sort of random sample. This has to be distinguished from the *conditional* variance, which is the variance of  $S_{t+1}$  given  $S_t$  and any other information available at time  $t$ . The conditional form may be thought of as the incremental variance, and the unconditional type as the total system variance.

Indeed, one important form of non-stationarity is where the stationary distribution, and hence the total system variance in the sense defined above, simply do not exist. We have already seen an example of this in the random walk, where the variance exploded over time, even though the conditional variance, which was the variance of the random increments  $\varepsilon$ , remained the same for all time. Lots of empirical processes behave this way, even if they do not have exactly the form of the random walk. Examples include stock prices, foreign exchange rates, and sometimes stock market returns. Such processes are said to have a *unit root*. The precise reference will not concern us too much here, but basically refers to the difference equation obeyed by the process, which would specify the way that  $S_t$  is derived in terms of  $S_{t-1}$  and past  $S$  values. There is a certain polynomial equation associated with this difference equation, and you can solve for its roots: if they are equal to or greater than one, then you know the whole process will explode over time, something like Figure 8.3 above.

Processes that explode in this way are often called *integrated processes*. To see why, recall the random walk. By a process of backward substitution, this can be written:

$$\begin{aligned} B_t &= \varepsilon_t + B_{t-1} \\ &= \varepsilon_t + \varepsilon_{t-1} + B_{t-2} \\ &= \varepsilon_t + \varepsilon_{t-1} + \varepsilon_{t-2} + \dots \quad \text{forever} \end{aligned}$$

The word “integrated” means *summed*, and you can see why this might be applied to the random walk. It can be viewed as the progressive outcome of all the summed past values of the innovations terms (recall these are the  $\varepsilon$ 's). In stationary processes, the influence of the innovations in the distant past dies away. For instance, consider the process defined by:



$$\begin{aligned}
y_t &= \lambda y_{t-1} + \varepsilon_t \\
&= \varepsilon_t + \lambda \varepsilon_{t-1} + \lambda^2 \varepsilon_{t-2} + \dots \\
&= \varepsilon_t + \lambda \varepsilon_{t-1} + \lambda^2 \varepsilon_{t-2} + \lambda^3 \varepsilon_{t-3} + \dots \text{ forever}
\end{aligned}$$

The second equation is obtained from the first by a process of backward substitution. If the parameter  $\lambda$  is less than one in absolute value, you can see that the influence of  $\varepsilon_{t-100}$  (i.e. the news 100 periods ago) is hit by a coefficient  $\lambda^{100}$ , which is very small. Compare this to the random walk where the coefficient of  $\varepsilon_{t-100}$  remains at one, just the same as the coefficient of  $\varepsilon_t$ . So whether  $|\lambda| < 1$  or  $|\lambda| \geq 1$  makes a great deal of difference. In the former, you can show that the unconditional variance  $\text{Var}(S_t)$  exists, but in the latter it does not exist – only the conditional variance  $\text{Var}(\varepsilon_t) = \sigma^2$ .

Suppose that we take the first differences of the random walk. We get:

$$B_t - B_{t-1} = \varepsilon_t$$

So the differenced random walk is just the innovations process. This is perfectly stationary; certainly the variance of  $\varepsilon_t$  exists, and in fact it is the same as  $\text{Var}(\varepsilon_{t-1})$  and so forth. We would say that the original  $B_t$  is  $I(1)$ , which is a shorthand way of saying that it is *integrated of order 1*, which, in turn, is a way of saying that if you difference it just once you will get down to a stationary process (denoted  $I(0)$ ). Sometimes processes are so badly non-stationary that you have to difference the difference to get down to a stationary process – these would be  $I(2)$  or worse. Most financial processes are  $I(1)$  or  $I(0)$ .

### Cointegrated processes

The above ideas extend to vector processes, in which we deal with two or more processes at a time, so that each can be regarded as an element of a vector. For instance, we could look at a vector process whose first element is the price of a specific stock, the second element is the price of the market, and the third is the value of the risk-free rate.

Often all the elements of a vector process are non-stationary, but there is an equilibrium relationship between them, which is itself stationary. For example, we saw in Chapter 7.5 that the CAPM specified that the return on a stock (number  $i$ , for instance) is related to the market return  $R_t$  and the risk-free rate  $\rho_t$  by:

$$r_{i,t} = \rho_t + \beta_i (R_t - \rho_t) + \varepsilon_{i,t}$$

Now it might be that  $r_{i,t}$  is non-stationary, and so are  $R_t$  and  $\rho_t$ . But the disturbance process  $\varepsilon_{i,t}$  is stationary. This is perfectly possible. It means that although the return on the stock, the market return, and the risk-free rate can drift all over the place, there is a kind of flexible rope always holding them together, so that the relationship is shocked (the  $\varepsilon$  terms) but only in a stationary way. In this case, the three variables  $r_{i,t}$ ,  $R_t$  and  $\rho_t$  are said to be *co-integrated*.

Where series are co-integrated, it turns out that you can always write a version of the relationship in which the change in one variable can be represented as the sum of the past changes in itself and the others, plus a term which represents the extent to which the equilibrium relationship was departed from in the past periods. This is the *error correction representation*; the error is in the equilibrium relationship, and the idea is that the variables adjust over time to remove these errors and keep it within co-ee of the equilibrium relationship.

Detailed study of these relationships is a little outside the scope of this book. In general, the purpose and value of time series models and methods is that they enable us to get a handle on the disturbance elements – the  $\varepsilon$  's. It is the latter – or, more specifically, their variances – that are so useful for pricing derivatives. The relationships are also useful in themselves for forecasting, in situations where you are after an uncovered gamble. In this type of scenario, it is the conditional mean – the  $\mu$  function – that now occupies center stage. It gives you the expected value of the stock price (or whatever physical you are interested in) as a function of the information available to you. When you model the conditional mean, you are in effect forming your own forecasting mechanism. The residual variance from this equation will give you some idea of the extent to which your forecast is likely to be wrong. In statistical jargon, the conditional variance can enable you to form confidence bands for your forecast. You would present the latter in the form of the point forecast – which is your computed value of the conditional mean – plus or minus one sigma for instance, which is the estimated standard deviation of the residual variance in your fitted relationship.

At this point, we draw a curtain on the discussion of stochastic processes and their applications. You already have a sense that this is an enormous topic, and an area of active research and applications programs both in academia and in industry. It has many important applications like derivatives pricing, forecasting, and risk management.

1. Can you think of some more examples - and perhaps better ones! – where you might want to hedge some price or return, but the available derivative instrument is not perfectly correlated with your target?
2. It has been said that the CAPM model and related portfolio models represent a form of hedge relationship. Suppose that you are a portfolio manager, and you find a stock that has a negative beta on the market.
  - (i) What does this mean about its expected return over the coming periods, relative to the risk-free rate?
  - (ii) Can you rationalize this in terms of the value of such a stock in terms of any specified hedging need?

Suppose, for instance, that the stock is an undertaking firm that specializes in expensive interments for highly strung stockbrokers. Would such a stock have a special value in your portfolio, and why?

3. (You may need a bit of help with this one.) Take a random walk process as specified in the form (2) of Section 8.4 in the text. You are now at time 0, and you know  $B_0$ .
  - (i) What is the predicted value of  $B_1$ ? What is the variance of the prediction, that is, the variance of the discrepancy between your predicted value and the actual outcome?
  - (ii) Do the same for  $B_2$  at time 2. Predict the value knowing only  $B_0$  at time 0, so you have to form the best predictor available and then find its variance, conditional on the information available at time 0.
  - (iii) Carry it on until some time  $T$  in the future. How does the predictive variance accumulate as the horizon  $T$  gets larger and larger?
4. Repeat Question 3 but this time with the process:

$$y_t = \lambda y_{t-1} + \varepsilon_t$$

Analyze the way in which the variance of your forecast increases or decreases with the horizon  $T$ , and relate this to the size of the absolute value of the parameter  $\lambda$ .

A vertical black bar with rounded ends. At the top, a large white number '8' is positioned. Below the number, the word 'Exercises' is written in white, bold, sans-serif font, oriented vertically.

## 8 Exercises

5. Use your data in Example 8.4 to explore the time value of the option. Graph the price of the SPI call option against maturity. Notice how the time value diminishes more rapidly, the closer you get to option maturity. (In hedging their portfolios, fund managers sometimes buy options on the SPI when they are worried about some forthcoming event, relying on keeping options costs down by selling the option in the near future before the time value starts decaying on them.)





# *Part Four*

**MANY VARIABLES**





---

# Nine

## Many Variables

Sometimes in finance, one has to handle situations where there are many variables to study or to control in some way. An obvious instance is in portfolio theory. In this context, having many stocks in your portfolio is regarded as good practice, as it enables you to diversify risk. But it does mean that one has to be able to manipulate not just one or two portfolio proportions, but a whole set of proportions. You can play around with the resulting decision problem by using elementary algebra, but the results can be atrocious and will almost certainly lead to error because things get messy pretty quickly. The methods of matrix algebra are a very natural way to handle this sort of situation. They can also help you remember results and properties a lot easier than you could otherwise! In this chapter we shall first give an abridged account of the basics of vectors and matrices, and then indicate the sort of situations where they can help. In particular, we will look into the formation of portfolios using mean—variance analysis.

## 9.1 Vectors and matrices

A vector is a one-dimensional array of numbers, usually written in the form of a column vector, which is a set of numbers placed one on top of the other. The array is given a name, and the name is usually written in bold type, as distinct from the names of ordinary variables, or scalars in this terminology, which are written in ordinary type. For example, we might have the vector:

$$\mathbf{x} = \begin{bmatrix} 2 \\ 3 \\ 5 \end{bmatrix} .$$

Vector  $\mathbf{x}$  is a 3 x 1 vector of numbers. We could denote the representative element of the vector  $\mathbf{x}$  as  $x_i$ , with the index  $i$  running from  $i = 1$  to  $i = 3$ :

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} ,$$

where in this case,  $x_1 = 2$ ,  $x_2 = 3$ ,  $x_3 = 5$ .

A matrix is a rectangular array of numbers. For example,

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 5 \\ 3 & 4 & 2 \end{bmatrix} .$$

Here the matrix  $A$  (we use uppercase but not bold) has order 3 x 3, meaning that it has three rows and three columns. This particular matrix is *square*; it has an equal number of rows and columns. The representative element of a matrix  $A$  is written  $a_{ij}$ , meaning the element in the  $i^{\text{th}}$  row and the  $j^{\text{th}}$  column.

Notice that we could regard the above matrix as being made up of three vectors, each of order 3x1, placed side by side:

$$A = [ \mathbf{a}_1 , \mathbf{a}_2 , \mathbf{a}_3 ]$$

$$\text{where: } \mathbf{a}_1 = \begin{bmatrix} 2 \\ 1 \\ 3 \end{bmatrix}$$

$$\mathbf{a}_2 = \begin{bmatrix} 1 \\ 0 \\ 4 \end{bmatrix}$$

$$\mathbf{a}_3 = \begin{bmatrix} 1 \\ 5 \\ 2 \end{bmatrix} .$$

Actually, to qualify as a matrix, it is not enough to be just a rectangular array of numbers. The array also has to obey certain operational laws, which stem from its mathematical background in coordinate geometry as the representation of a linear operator. For our purposes, the most important of these rules concerns multiplication.

Suppose you have a vector like  $\mathbf{x}$  above. Then the product  $A\mathbf{x}$  has the following meaning: it is a vector whose elements are obtained by means of the scheme depicted below.

$$\begin{array}{c} \downarrow \\ \rightarrow \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} ? \\ ?? \\ ??? \end{bmatrix} \text{ with } ?? = a_{21}x_1 + a_{22}x_2 + a_{23}x_3 . \end{array}$$

In this scheme, to get the second element of the product  $A\mathbf{x}$ , we go across the second row of  $A$  and multiply each element of that row with a succeeding element of the vector  $\mathbf{x}$ , adding the results. We do the same for the first and third elements of the product. Now try this out by taking the product  $\mathbf{y} = A\mathbf{x}$ , where  $\mathbf{x}$  and  $A$  are the earlier numerical examples. Verify that:

$$\mathbf{y} = \begin{bmatrix} 12 \\ 27 \\ 28 \end{bmatrix} .$$

You will see that this process will only work if the number of columns of  $A$  is equal to the number of elements (rows, in this case) of the vector  $\mathbf{x}$ . In this case,  $A$  and  $\mathbf{x}$  are said to be *conformable*.

Example: The *vector of ones*, or summation vector, is particularly useful:

$$\mathbf{1} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

Can you verify that the vector  $A\mathbf{1}$  is equal to the sum of the columns of  $A$ ?

Now suppose that we have another matrix  $X$ , with columns  $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$ . Each of the columns can be treated as a column vector itself. So we can attach a meaning to the product  $Z = AX$ . Each column  $\mathbf{z}$  of the matrix  $Z$  will be obtained as the product of  $A$  times the corresponding column of the matrix  $X$ , the latter treated as a matrix - vector product, in the manner just outlined. In other words,

$$Z = AX = [A\mathbf{x}_1, A\mathbf{x}_2, A\mathbf{x}_3, A\mathbf{x}_4]$$

Again, this will only work if the number of columns of  $A$  is equal to the number of rows of  $X$ , in which case,  $A$  and  $X$  are conformable. If this is the case, we can build up the product of two matrices by taking the first times each column of the second in turn, and putting them side-by-side. The order of the resulting product is equal to the number of rows of  $A$  by the number of columns of  $X$ . A good mnemonic is to write the orders of  $AX$  as  $[3 \times 3] \times [3 \times 4]$ . The two inner numbers have to be the same for conformability, and the two outer numbers give the order of the result, as  $[3 \times 4]$ .

You will notice that you have to be a bit careful about the order in which you write the product. Thus  $XA$  means something quite different from  $AX$ . Indeed, even if the first product does exist, it does not follow that the second product exists - the conformability requirements are quite different.

In Question 1 at the end of this chapter we have given several matrix multiplication problems, using easy numbers. If you are not already conversant with matrix products, you should work through these before proceeding.

The *transpose* of a matrix  $A$ , written  $A'$ , is the matrix constructed by turning the rows of  $A$  into columns. For example:

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 5 \\ 3 & 4 & 2 \end{bmatrix} \quad \text{then} \quad A' = \begin{bmatrix} 2 & 1 & 3 \\ 1 & 0 & 4 \\ 1 & 5 & 2 \end{bmatrix} .$$

A square matrix,  $A$  is symmetric if  $a_{ij} = a_{ji}$ . In other words, if you interchanged rows and columns, the new matrix would be just the same as the old one. You will notice that if  $A$  happens to be symmetric, then  $A = A'$ . Of course,  $A$  can only be symmetric if it is square to begin with.

You can have the transpose of a vector, too. In this case, a single column becomes a single row, often called a *row vector*.

$$\mathbf{x} = \begin{bmatrix} 2 \\ 3 \\ 5 \end{bmatrix}$$

$$\mathbf{x}' = [ 2 \quad 3 \quad 5 ]$$

Now suppose we take the matrix,  $A$ , which is symmetric:

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 4 \\ 1 & 4 & 2 \end{bmatrix} .$$

(Check that  $A$  is symmetric, in terms of the definition given above.)

Now we form the following triple product:

$$q = \mathbf{x}' A \mathbf{x} .$$

Let us check that we can do this properly to begin with, that is, that everything is conformable. Take the last product,  $A\mathbf{x}$ .  $A$  has 3 columns and  $\mathbf{x}$  has three rows, so the product will exist, and be of order  $3 \times 1$ , the number of rows of  $A$  by the

number of columns of  $\mathbf{x}$  (which, incidentally, is one, since it is a vector). The result is a vector of order  $3 \times 1$ . The transpose  $\mathbf{x}'$  is a row vector of order  $1 \times 3$ . The last step is to multiply a row vector by column vector:  $1 \times 3$  by  $3 \times 1$ , and you can see that the result exists and is of order  $1 \times 1$ . This means that it is just a single number, or a *scalar*. Go ahead and try it. You will find that the result is:

$$q = 2x_1^2 + 2x_1x_2 + 2x_1x_3 + 8x_2x_3 + 2x_3^2$$

Could you have guessed what the coefficients of this expression are, just by looking at the matrix A?

This sort of expression is called a *quadratic form*. Every term in it is of power two: either a square like  $x_1^2$  or a product like  $x_1x_2$ . Conversely, every quadratic form can always be represented as  $\mathbf{x}' \mathbf{A} \mathbf{x}$ , for some symmetric matrix A.

## 9.2 Matrix inverses and equation solving

Suppose we have an equation

$$3x = 6$$

We assume you are able to solve this, hopefully without much difficulty. We could write the solution in the form:

$$x = 3^{-1} 6$$

where you will recall that  $3^{-1}$  is just another way of writing  $1/3$ . The answer, of course, is  $x = 2$ .

Now we are going to do the same sort of thing for solving systems of equations, which are linear simultaneous equations such as :

$$3x_1 + 2x_2 = 4$$

$$x_1 + 5x_2 = 1$$

You may know how to work out the solutions for  $x_1$  and  $x_2$  already. It's easy for this  $2 \times 2$  set of equations and variables, but you'll need something more than high school algebra when it comes to solving  $14 \times 14$  problems!

Before we can move on, there is a little more work to do on matrices. One special matrix will be important: the identity matrix, denoted I. This is a matrix with ones

all the way down the leading diagonal—which is the diagonal leading from top left to bottom right—and zeros in all the off diagonal positions. For example the 2 x 2 identity matrix is:

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Write down the 3x3 identity matrix.

The identity matrix plays the role of the number one (1) in matrix work. If you take any conformable matrix and multiply it by I, you get the same matrix you started with; thus  $IA = A$  or  $AI = A$ . Similarly, if  $\mathbf{x}$  is a vector,  $I\mathbf{x} = \mathbf{x}$ . (Would  $\mathbf{x}I$  make sense?) Now take a matrix A. It has to be square. Under suitable circumstances—more on this later—we can find another matrix, denoted  $A^{-1}$ , such that:

$$A^{-1}A = I$$

If we can find such a matrix, it is said to be the inverse of A. For us to be able to find such a matrix, we need the matrix A to be of full rank. By this we mean that the columns of A must be genuinely different. For example, consider the matrix:

$$A = \begin{bmatrix} 2 & 1 & 2 \\ 1 & 0 & 0 \\ 3 & 4 & 8 \end{bmatrix},$$

or the matrix

$$B = \begin{bmatrix} 2 & 4 & 1 \\ 1 & 11 & 5 \\ 3 & 7 & 2 \end{bmatrix}.$$

You can see why matrix A might have a problem. The third column of A is just twice the second, so in a sense the columns are not really different. The case B is a bit harder, but you might be able to see that the second column is equal to the first plus twice the third. So once again, the columns are not genuinely different, since one of them can be manufactured out of the others. This property is called *linear dependence*. To be able to compute its inverse, we must have all the columns of

A as linearly independent, meaning that you cannot get one from the others by simply taking some linear combination of them. Such a matrix is said to be of *full rank* or *non-singular*. In practice, the computer will tell you when the matrix you are trying to invert is singular, so you don't have to bother. The computer will also do the inverting, that is, finding  $A^{-1}$ .

Now we move on to the use of matrix inverses. Suppose we have the following set of three simultaneous equations with three unknowns  $x_1, x_2, x_3$ :

$$2x_1 + x_2 + x_3 = 2$$

$$x_1 + 5x_3 = 3$$

$$3x_1 + 4x_2 + 2x_3 = 5$$

Using your newfound skills in writing things in matrix-vector form, verify that the above can be written as:

$$\mathbf{A} \mathbf{x} = \mathbf{b} , \tag{1}$$

where  $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$

$$\mathbf{b} = \begin{bmatrix} 2 \\ 3 \\ 5 \end{bmatrix}$$

and  $A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 5 \\ 3 & 4 & 2 \end{bmatrix}$ .

You can see that all this is quite general. Any set of linear equations can always be written in the form (1).

Now multiply both sides of Equation (1) by the inverse  $A^{-1}$ . (We don't mean to



calculate  $A^{-1}$  numerically; this is just matrix algebra we are doing here). More precisely, multiply everything on the left hand side, or *pre-multiply* by  $A$ . Since  $A^{-1}A = I$ , the identity matrix, we get:

$$\mathbf{x} = A^{-1} \mathbf{b}$$

In effect, this has solved the system, just as  $x = 3^{-1} 6$  solved the simple one-variable equation. The computer can work out  $A^{-1}$  and multiply the given right hand side vector by the inverse, and—*Bingo!*—there it is: the solutions for  $x_1, x_2, x_3$ .

Actually, the inverse of the matrix  $A$  is:

$$A^{-1} = \frac{1}{23} \begin{bmatrix} 20 & -2 & -5 \\ -13 & -1 & 9 \\ -4 & 5 & 1 \end{bmatrix}$$

Can you check that  $A^{-1}A = I$ ? Now find the solution for  $\mathbf{x}$ . We will not worry here about the algorithms used to calculate the inverse. Basic versions can be found in every textbook of linear or matrix algebra.

### 9.3 Statistics with matrices

A lot of useful financial econometrics is concerned with the mutual relationships between many variables, such as a set of securities. Suppose that  $x_1, x_2, \dots, x_n$  is a set of  $n$  random variables. Note that  $n$  does not mean the sample size here - it is the number of random variables. For instance, in the data you have  $n = 14$  individual stocks. We can combine the random variables into a random vector  $\mathbf{x}$  which will be of order  $n$ :

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \cdot \\ \cdot \\ x_n \end{bmatrix}$$

The vector random variable  $\mathbf{x}$  will have a mean  $\boldsymbol{\mu}$ :

$$E \mathbf{x} = \begin{bmatrix} E x_1 \\ E x_2 \\ \cdot \\ \cdot \\ E x_n \end{bmatrix} = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \cdot \\ \cdot \\ \mu_n \end{bmatrix} = \boldsymbol{\mu} .$$

Moreover, there will be variances and covariances between the various random variables. For instance,  $\sigma_1^2 = E(x_1 - \mu_1)^2$ , or  $\sigma_{12} = E(x_1 - \mu_1)(x_2 - \mu_2)$ . We can assemble all the variances and covariances into a matrix:

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \sigma_{12} & \cdot & \cdot & \sigma_{1n} \\ \sigma_{12} & \sigma_2^2 & \cdot & \cdot & \sigma_{2n} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \sigma_{1n} & \cdot & \cdot & \cdot & \sigma_n^2 \end{bmatrix}$$

Strictly speaking, we should write the (2,1)<sup>th</sup> element as  $\sigma_{21}$  and not  $\sigma_{12}$ , but they are equal so we just use one and not both. A similar convention applies to all the other off-diagonal elements. We also use  $\sigma_1^2$  instead of  $\sigma_{11}$  (etc.) for the diagonal elements.

Note that the  $\Sigma$  symbol does not represent the summation sign in this context. It is simply the uppercase \* (for matrices, remember) version of the lower case  $\sigma_{ij}$  inside. You will notice that it is a symmetric matrix, simply from the way it is constructed.

Often the random vector  $\mathbf{x}$  is specified to have a given probability distribution. An example is multivariate Normal. This has a similar shape to the unidimensional Normal, but now has to be plotted in higher dimensions. For example, if you had normal density for two random variables, this would look like a bell in 3-D. We would say that  $\mathbf{x}$  has the multivariate  $N(\boldsymbol{\mu}, \Sigma)$  distribution function or density.

Another way of writing the covariance matrix  $\Sigma$  is as follows:

$$\Sigma = E(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})'$$

You can try writing this out in full by multiplying  $(\mathbf{x} - \boldsymbol{\mu})$  by  $(\mathbf{x} - \boldsymbol{\mu})'$ , and then taking

the expectations of each of the elements of the resulting matrix. Remember that the prime means transpose, so the orders are  $[n \times 1] \times [1 \times n]$ , and the result will be  $[n \times n]$ .

#### 9.4 Portfolio theory<sup>1</sup> with matrices

At last we are at the point where we can apply these concepts. Take a set of  $n$  securities, and call their returns  $r_1, r_2, \dots, r_n$ . Each of these securities will have a conditional expectation, given the information available to the investor. This may depend on time, in the sense that the returns are to accrue over the interval from  $t$  to  $t + 1$ . At instant  $t$ , the investor can form his or her expectations of the expected return over the current period: for the representative security  $i$ , for instance:

$$E r_{i,t} = \mu_{i,t}^*$$

For example, investors may be operating in terms of some sort of CAPM model. They know that returns this period (between instants  $t$  and  $t + 1$ ) will certainly depend on the risk free rate ( $\rho_t$ ) to hold for investment over the coming period, which is of course, known in advance and therefore part of the investor's information set. Hence we might specify:

$$\mu_{i,t}^* = \rho_t + \mu_{i,t} \quad (2)$$

where the  $\mu_{i,t}$  might be formulated in terms of what the investor thinks the market as a whole will do (the *market risk premium*), plus some version of how he or she thinks the individual stock is related to the market. We will leave open the precise way that investors are imagined to estimate their means for the coming period.

You can see that if you assume some sort of decomposition like (2), you might as well simply redefine your returns  $r_{i,t}$  as the *excess returns* relative to the risk-free rate, so this is what we shall do, both in the theory and in the practical work.

*NOTE: In the following discussion, all returns and return symbols refer to excess returns, that is, the original return minus the risk-free rate.*

<sup>1</sup>The reader should note that there are various versions of portfolio formation in practice, which may not be the same as the ones discussed here. A common procedure is to use the unconditional means and variances, often computed over a very long period, to calculate the efficient frontier, which means the set of portfolios that maximize the average return for any given variance. The investor is then asked (in effect) to choose a portfolio along this frontier, in accordance with his or her preferred risk attitudes. In such calculations, one would use the *actual* returns, rather than the excess returns which we used here.

As noted, the excess returns  $r_{i,t}$  will have a mean  $\mu_{i,t}$  which may depend on time. However, in the interest of notational simplicity, we shall simply suppress the time dependence, and write the means of the excess returns  $r_i$  as simply  $\mu_i$ .

Likewise, the excess returns will have a mutual covariance matrix as perceived by investors, for the returns over the coming period. We shall denote this by  $\Sigma$ . Again this may depend on time, via the investor's information set. For example, the investor may perceive this coming period as a time of greater volatility than usual - a substantial oil shock has just occurred, and times are suddenly rather uncertain.

In summary, the excess returns vector  $\mathbf{r}$  is assumed to have some sort of distribution over the coming period, with mean  $\boldsymbol{\mu}$  and covariance matrix  $\Sigma$ . For the moment, we shall assume that these are known. The task then is to find out, in light of this knowledge, how much of each security should be held in the investor's portfolio. To answer this question, imagine that you have \$1 available for investment, and you have to parcel this out among the set of  $n$  risky stocks, bonds or whatever the securities actually are. You will be after the portfolio proportions  $x_1, x_2 \dots x_n$ . Note that these proportions must sum to one. Collect these proportions into a vector  $\mathbf{x}$ .

Once you do this, the *portfolio (excess) return* is itself going to be a random variable given by:

$$p = x_1 r_1 + x_2 r_2 + \dots + x_n r_n = \mathbf{x}' \mathbf{r}$$

You can see the mean portfolio return (henceforth, "return" refers to excess return) is given by:

$$E p = \mu_p = x_1 \mu_1 + x_2 \mu_2 + \dots + x_n \mu_n = \mathbf{x}' \boldsymbol{\mu}$$

Likewise, as a random variable, the return has a variance given by:

$$\sigma_p^2 = E (p - \mu_p)^2$$

Using your newfound expertise in matrix-vector manipulation (though this one may take some experience!), try to show that the portfolio variance can be written as:

$$\sigma_p^2 = \mathbf{x}' \Sigma \mathbf{x} . \quad (3)$$

<sup>2</sup> Do not confuse this vector with the  $\mathbf{x}$  of Section 9.3, which was a general description of a random vector. Here we follow standard notation in portfolio theory and reserve  $\mathbf{x}$  for the non-random portfolio proportions, and it is  $\mathbf{r}$  that is the random vector.

In other words, the portfolio variance is just a quadratic form in the portfolio proportions  $\mathbf{x}$  with elements of the covariance matrix  $\Sigma$  forming the coefficients.

### Mean - variance portfolios

How should you choose the right portfolio proportions—that is, the vector  $\mathbf{x}$ ? To do this you have to have some objective function. Possible choices are:

#### A. *Minimum variance*

Choose  $\mathbf{x}$  to minimize the portfolio variance (3) subject also to the sum of the  $x$  being one. This result is the *minimum variance portfolio*. It can be shown that the proportions of the minimum variance portfolio are given by:

$$\mathbf{x} = \left( \frac{1}{\mathbf{1}'\Sigma^{-1}\mathbf{1}} \right) \Sigma^{-1}\mathbf{1} \quad (4)$$

Recall that the vector  $\mathbf{1}$  is a vector with all its elements just 1, and it represents the process of summation. Therefore in the top of (4), the product  $\Sigma^{-1}\mathbf{1}$  turns out to be a vector which is equal to the sum of the columns of the inverse of the covariance matrix. The denominator of the right hand side of (4) is a normalizing scalar which is equal to the sum of all the elements of the inverse of the covariance matrix. It makes all the elements of  $\mathbf{x}$  sum to one.

Thus, the minimum variance portfolio gives weights to each security,  $I$ , proportional to the sum of the  $i^{\text{th}}$  row of the elements of  $\Sigma^{-1}$ . Note that we have defined the minimum variance portfolio in terms of excess returns, rather than nominal returns, though it is not uncommon to use nominal returns.

#### B. *The CAPM portfolio*

This turns out to be the portfolio that would be held by investors who possess a utility function for wealth that says that they prefer more money to less, but that the marginal utility of the extra unit of wealth is declining. Such investors display *risk-averse* behavior. Under suitable assumptions (which we won't go into here), this generates a CAPM type of capital market equilibrium, in which everybody holds the same portfolio proportions, given by:

$$\mathbf{x} = \left( \frac{1}{\mathbf{1}'\Sigma^{-1}\boldsymbol{\mu}} \right) \Sigma^{-1}\boldsymbol{\mu} \quad (5)$$

You can see that this is the same form as the minimum variance portfolio with the mean of the excess returns, namely  $\boldsymbol{\mu}$ , replacing the summation vector  $\mathbf{1}$ . Note that  $\boldsymbol{\mu}$  has to be the mean of the *excess* returns. The interpretation is that the weights are proportional to the vector  $\Sigma^{-1} \boldsymbol{\mu}$ , which are subsequently adjusted to sum to one.

To compute the optimum portfolio, of either type, we need to know the covariance matrix of excess returns, conditional on current information  $I_t$ . Once we know this, the computer can invert and calculate the portfolio proportions for us.

Finally, with the portfolio proportion  $\mathbf{x}$  solved for both A and B, you can form the mean and variance for the portfolios:

$$\begin{aligned}\mu_p &= \mathbf{x}'\boldsymbol{\mu} && \text{portfolio mean;} \\ \sigma_p^2 &= \mathbf{x}'\Sigma\mathbf{x} && \text{portfolio variance}\end{aligned}$$

## 9.5 Practicum

Now we will use the historical data on US stocks to compute the portfolios. The essential first steps are to estimate the conditional mean and covariance matrix for the coming period. One simple assumption that is often made is that the conditional means for excess returns in each period are exactly the same. This virtually amounts to assuming that the historical excess returns can be treated as though they were simply a sample from some kind of stationary distribution. So in that case, we could use for  $\boldsymbol{\mu}$  simply the sample mean of the excess returns over the observation period, and for the  $\Sigma$ , the sample covariances of excess returns over the observation period. More advanced treatments try not to make such strong assumptions, allowing both the means and the covariances to change over time, but to get the ball rolling we will make the simpler assumption here.

The first step is to estimate the sample mean of each of the 14 stocks on the disk. The next is to estimate the variances and covariances. Since there are 105 different variance/covariances, we'll take pity on you and suggest that you choose just five stocks to form the portfolios. For our example, we have simply chosen a variety of stocks to provide a nice industrial coverage. We stress that the selection does not imply any judgement whatever about the quality of the stocks. The five stocks chosen to illustrate things are:

IBM  
GM  
MMM  
Marine Drill  
Wendy's

Return to your Excel worksheet where you have stored the excess returns data, that is, the data from which the effective 30-day CD has been subtracted.

Now find a clear section of the worksheet in which to perform the calculations. In ours, the area just below the data starting in cell [A89] was chosen as it was empty right over to column S for all cells south of the data. For ease of explanation, we will continue the description of setting up the series of calculations as though you have chosen the same area. We are going to construct a calculation block consisting of the mean and variance of the excess return for each stock, followed by a matrix of the covariances for each pair of stocks. Completed, the block should appear as it does in Table 9.1 below.

Starting in cell [A90] type the word **MEAN**; type **VAR** in cell [A91] and then **COVAR** in cell [A92]. Into cell [B92], next to "COVAR" (and this is important, as we shall explain later), type the name for your matrix. This is best given as a single letter, such as **A**, for the ease of performing matrix operations.

Then type the names of the stocks in the above order (as they are roughly lined up in your data) across the top of the area designated for your calculation block--for us, cells [B89] through [F89]. Input the same names, again in alphabetical order down column A, into cells [A93] through [A97].

Beginning in cell [B90], next to the word "MEAN," find the mean of the excess returns for IBM. This is accomplished by typing the formula:

$$=AVERAGE(\textit{cell range for IBM excess returns})$$

On our worksheet:

$$=AVERAGE(B2:B60)$$

You can then copy and paste the formula into the same row in the four cells to the right, directly under the other stock names. You will need to alter the cell range reference in each formula so that it points to the excess returns of the correct

stock. For example, the formula to calculate the mean returns for MMM will be *AVERAGE(F2:F60)*; for Wendy's the formula will read *AVERAGE(O2:O60)*.

Under the mean (or "AVERAGE"), use Excel to calculate the variance of excess returns for the chosen five stocks. Either type or use the Function Wizard to put the following formula into the cell directly below the mean for IBM:

*=VAR(cell range for IBM excess returns)*

On our worksheet:

*=VAR(B2:B60)*

Then copy and paste the formula into the four cells to the left in the same row, again adjusting the cell references to line up with the correct returns data. You should now have completed the means and variances.

Now calculate the covariances. Run the cell pointer down the IBM column to the point where it intersects with the GM row, cell [B94] on our worksheet. Using the Paste function button (labeled *fx*) on the standard toolbar, select *Statistical* as the Function Category, *COVAR* as the Function Name and click on *[OK]*.

For the next dialog box, respond as follows:

Array 1:    *\$C\$2:\$C\$60*    to select the GM Data, then hit the tab key

Array 2:    *B\$2:B\$60*      to select the IBM data

Then click *[OK]*.

Note that the full cell references for Array 1 are absolute, whereas the cell references for Array 2 are relative with respect to column, but absolute with respect to row. This will simplify the copy and paste actions which follow.

Now, copy the contents of this first GM cell across the row into cells [D94], [E94] and [F94] where the GM row intersects the MMM, Marine Drill and Wendy's columns. Again, you must adjust the second cell reference in each formula to point to the correct stock returns. For example, the formula for the covariance of GM with Marine Drill reads *=COVAR (\$C\$2:\$C\$60,N\$2:N\$60)*.



What you are building is a *variance-covariance matrix*. For all the cells which represent the intersection of one stock with different stocks, you will calculate the co-variance between the two intersecting stocks. For the time being, the diagonal of the matrix, where the row for a given stock intersects with the column for the same stock, should be left blank. Later, these cells will be occupied by the variances for the respective stocks. The remaining covariance cells can be filled using cut and paste, and editing where necessary so that the two arrays referenced by the formula in each of the covariance cells contain the excess returns data for the two "intersecting" stocks.

Finally, put the variances for the individual stocks into the diagonal cells running from the top left to the bottom right of the matrix. We found that this was easily done by typing:

```
=B94
    =C94
        =D94
            =E94
                =F94
                    diagonally across the matrix.
```

That is, type an equals sign and the cell reference of the variance you have already calculated for each stock. Completed, the variance-covariance matrix, and accompanying summary statistics should appear as in Table 9.1 below.

Problem 9.4 at the conclusion of this chapter gives you an algorithm for constructing the covariance matrix by using matrix operations.

**Table 9.1 Covariance matrix for stocks**

	IBM	GM	MMM	Marine Drl	Wendys
MEAN	0.023063	0.007322	0.008952	0.039054	0.00502
VAR	0.010548	0.009299	0.005281	0.04319	0.008576
COVAR	A				
IBM	0.010548	0.003366	0.000867	0.003023	0.001137
GM	0.003366	0.009299	0.001086	0.003183	0.001047
MMM	0.000867	0.001086	0.005281	0.004897	0.00058
Marine Drl	0.003023	0.003183	0.004897	0.04319	0.002159
Wendys	0.001137	0.001047	0.00058	0.002159	0.008576

### General instructions for Excel matrix operations

When you construct your matrix in Excel, give it a name and make a note of what it contains. The name *must* be placed in the cell directly above the top left hand corner of the matrix, which we have already done for the current data, putting the name A in cell [B92].

To perform matrix operations using the data you have set up, it is now necessary for you to formally notify the Excel program that the data you will be using exists in a named block of cells in your worksheet. You can do this as follows:

1. Using the mouse pointer, highlight the block of cells holding the input matrix. For purposes of the current application, this is the variance-covariance matrix, that is, cells *B93:F97* if you followed our set up.
2. Run the mouse pointer up to the menu bar and click *Insert, Name, Define...*
3. The *Define Name* dialog box will appear. In the *Names in Workbook* input box, there should appear the letter A, the name you provided for your matrix in cell [B92]. (If not, type in the name of your matrix.)

The *Refers to:* input box should list the range of your matrix. If it doesn't, click in the box and edit using the delete and character keys to list the correct cells. For the present application this is *B93:F97*. Click [*OK*].

Now we are about ready to run. There are at least two methods for performing matrix operations in Excel. The first method requires the least input time, but the second method might be easier for those of you who find the final keystroke combination difficult.

If you get through Method 1, skip Method 2 and go directly to the following section on Mean-variance portfolios: solutions.

#### Matrix operations: Method 1

1. For the desired matrix operation, you must determine the order (i.e., dimensions) of the resulting matrix, and highlight the cells that it will occupy on your spreadsheet.
2. With the "result" cells highlighted, type an equals (=) sign followed by the required matrix command. (This could be, for example: `=MINVERSE(A)`)

*Don't* hit the [Enter] key at this point.

- The simplest way to obtain the result matrix now is to strike three keys at once:

[Ctrl] + [Shift] + [Enter]

(Press [Ctrl] and [Shift] using your left hand, and while they are held down, then press [Enter]).

### Matrix operations: Method 2

- As with Method 1, for the desired matrix operation you must determine the order (i.e. dimensions) of the resulting matrix, and reserve those cells on your spreadsheet.
- In the row directly above your result matrix, type the number one (1) in the far left cell, followed by two (2) in the next column, and so on for the required number of columns. (These cells provide the column references).
- In the column directly to the right of your result matrix location, type the number one (1) in the first row, two (2) in the second, and so forth down to the final row. (These cells will provide the row references).
- In the top left cell of the result matrix, type your formula in the following format:

*=INDEX(name of the Excel matrix operation [e.g. MINVERSE] (name of the input matrix [e.g. A, as we have chosen]), row reference, column reference).*

Suppose, for example, that your result matrix is to rest in cells *B123:F127*. One possible layout is shown below, with cell contents designated in bold and cell names in italics.

<i>A122</i>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>
<b>1</b>	<i>B123</i>	<i>C123</i>	<i>D123</i>	<i>E123</i>	<i>F123</i>
<b>2</b>	<i>B124</i>	<i>C124</i>	<i>D124</i>		
<b>3</b>	<i>B125</i>		<i>D125</i>		
<b>4</b>	<i>B126</i>				
<b>5</b>	<i>B127</i>				<i>F127</i>

To obtain the result of the matrix operation for cell B123, into B123 we would type the formula:

`=INDEX(Matrix operation(matrix name),$A123,B$122)`

If this command is copied into cell [B124], it will automatically change to:

`=INDEX(Matrix operation(matrix name),$A124,B$122),`

which is exactly what we want. Pasted into [C123], it will give:

`=INDEX(Matrix operation(matrix name),$A123,C$122),` correctly again.

In fact, the formula will provide the correct row and column references if cut and pasted throughout the result matrix. It is a bit of an effort to set up, but in doing so, you have learned something quite useful about cell referencing, in addition to avoiding the three-at-a-time keystroke combination.

Performing a matrix operation without the three-at-once keystroke combination will generate an answer which includes only the top left hand element for the result matrix from the operation. The Index command provides a way to access the remaining elements of the result matrix.

### Mean-variance portfolios: solutions

Having set up the covariance matrix, the next task is to generate the mean-variance portfolios of Section 9.4. We'll use Method 1 to perform matrix operations, but you can apply Method 2 just as easily.

A. Starting with the Minimum Variance portfolio, we first have to find the inverse of the covariance matrix,  $\Sigma^{-1}$ . Begin by highlighting a 5 x 5 block of cells a few cells below the original covariance matrix. (We used cells [B100] to [F104].) Type:

`= MINVERSE(A)` into the top left hand cell of the highlighted block.

Then press [Ctrl] + [Shift] + [Enter].

With a bit of luck, you will have created the new matrix,  $\Sigma^{-1}$ , which appears as Table 9.2.

**Table 9.2**  $\Sigma^{-1}$ 

108.9531	-36.464	-5.87672	-3.841636	-8.63143
-36.464	124.941	-14.5443	-4.592299	-8.2745
-5.87672	-14.5443	214.937	-22.57479	-6.2945
-3.84164	-4.5923	-22.5748	26.5248	-4.08114
-8.63143	-8.2745	-6.2945	-4.081138	120.2174

Referring back to Equation (4), we must now multiply  $\Sigma^{-1}$  by the vector  $\mathbf{1}$ . The effect of this multiplication is the same as adding across the columns of  $\Sigma^{-1}$ , row by row. We can do this very easily on the worksheet. Check that you obtain the vector shown as Table 9.3.

**Table 9.3**  $\Sigma^{-1}\mathbf{1}$ 

54.13937
61.06588
165.6467
-8.56506
92.93585

By summing the five elements of the vector  $\Sigma^{-1}\mathbf{1}$ , we obtain the scalar (ordinary number):

$$\mathbf{1}'\Sigma^{-1}\mathbf{1} = 365.2227$$

The vector of weights for the five securities is then obtained via Equation (4) as:

$$\mathbf{x} = \left( \frac{1}{\mathbf{1}'\Sigma^{-1}\mathbf{1}} \right) \Sigma^{-1}\mathbf{1} ,$$

that is, by multiplying the vector  $\Sigma^{-1} \mathbf{1}$ , by the inverse of the scalar above.  
This is easily done on the spreadsheet to give Table 9.4:

**Table 9.4**  $(1 / \mathbf{1}'\Sigma^{-1} \mathbf{1}) * \Sigma^{-1} \mathbf{1}$

0.148237
0.167202
0.45355
-0.02345
0.254464

Clearly, the Minimum Variance portfolio is heavily weighted towards the MMM stock and Wendy's - by just over 70%. The minus sign next to the fourth weight indicates shorting a small amount of the Marine Drill stock.

- B. For the CAPM portfolio, we again use the inverse of the covariance matrix  $\Sigma^{-1}$ , that we produced for the first portfolio above.

Following Equation (5), we must next multiply by  $\boldsymbol{\mu}$ , the vector of the stock mean (excess) returns. You must first create the vector as a column of five elements which we did as:

=B90  
=C90  
=D90  
=E90  
=F90

These are the cell references for our original calculation of the return means.

This produces the required vector:

$$\boldsymbol{\mu} = \begin{bmatrix} .023063 \\ .007322 \\ .008952 \\ .039054 \\ .00502 \end{bmatrix} .$$

Use the Excel function MMULT (matrix multiplication) to compute  $\Sigma^{-1}\boldsymbol{\mu}$ . To perform this function, first highlight a column of five empty cells for the result vector. Into the top cell type:

$$= \text{MMULT} (\text{cell range for } \Sigma^{-1}, \text{cell range for stock means}),$$

which on our worksheet was:

$$= \text{MMULT}(\$B\$100:\$F\$104,\$H\$90:\$H\$94).$$

Then strike the combination: [Ctrl] + [ Shift] + [Enter], to produce:

$$\Sigma^{-1}\boldsymbol{\mu} = \begin{bmatrix} 1.999821 \\ -0.277177 \\ 0.768771 \\ 0.691117 \\ 0.128052 \end{bmatrix} .$$

The front end of the right hand side of Equation (5),  $\mathbf{1}'\Sigma^{-1}\boldsymbol{\mu}$ , is just the sum of the elements of the vector  $\Sigma^{-1}\boldsymbol{\mu}$ , that is, the scalar which, for the data selected, comes to:

$$\mathbf{1}'\Sigma^{-1}\boldsymbol{\mu} = 3.310584$$

The vector  $\mathbf{x}$  of the weights for the selected securities to the CAPM portfolio is then:

$$\mathbf{x} = \left( \frac{1}{\mathbf{1}'\Sigma^{-1}\boldsymbol{\mu}} \right) \Sigma^{-1}\boldsymbol{\mu} .$$

Its solution is given by:

$$\mathbf{x} = \begin{bmatrix} 0.604069 \\ -0.083724 \\ 0.232216 \\ 0.20876 \\ 0.03868 \end{bmatrix} .$$

This time the portfolio weights indicate a high proportion of our capital in IBM shares (60.4%), while we are directed to short sell a small amount of GM (amounting to 8.4% of our capital amount).

Finally, we will actually calculate the mean excess return of the two portfolios, again on an historical basis, and also their standard deviations. The standard deviation of the portfolio is sometimes called the capital at risk in portfolio management literature.

The variance of each portfolio excess return is calculated as:

$$\sigma_p^2 = \mathbf{x}' \Sigma \mathbf{x} \quad ,$$

while the mean excess return is computed as:

$$\mu_p = \mathbf{x}' \boldsymbol{\mu} \quad .$$

The results of these computations for the two portfolios are shown as Table 9.5 below. As you can see, the Minimum Variance portfolio (A) ended up with a standard deviation of about 62% of the CAPM portfolio (B), while its mean excess return was less than one-half of that achieved with portfolio (B).

**Table 9.5**  $\sigma_p$  ,  $\mu_p$

Portfolio	$\sigma_p^2$	$\sigma_p$	$\mu_p$
(A) Minimum Variance	.002738	.052326	.009065
(B) CAPM	.007172	.084689	.023744



1. (i) Determine which of the following matrix products can be found, and find them.

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 5 \\ 3 & 4 & 2 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 2 \\ 1 & 0 \\ 3 & 4 \end{bmatrix}.$$

- (a) AB    (b) AC    (c) CA    (d) CB    (e) BC

- (ii) Suppose that you have an  $N \times 1$  vector  $\mathbf{1}$ , whose every element consists of the number one (1). Now take a set of sample observations on some random variable  $x$ , namely  $x_1, x_2, \dots, x_N$ , and assemble them into a column vector  $\mathbf{x}$ . What does the product  $\mathbf{1}'\mathbf{x}$  look like? The vector  $\mathbf{1}$  is called the *summation vector* – it does the same thing as the  $\Sigma$  sign.
2. In Section 9.2 you will find an answer to the inverse of the matrix below. Check that answer is correct using Excel (See Section 9.5).

$$A = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 5 \\ 3 & 4 & 2 \end{bmatrix}$$

3. Consider the following set of equations:

$$2x_1 + x_2 + 2x_3 = 1$$

$$x_1 = 2$$

$$3x_1 + 4x_2 + 8x_3 = 1$$

- (i) From first principles, try to solve it. Does this set have a unique solution?
- (ii) Write the set of equations in matrix-vector form.
- (iii) Relate the solution problem to the linear dependence of the system matrix.

4. Here are some useful manipulations with statistical applications.
- (i) Suppose that you have a  $N \times k$  matrix of  $N$  observations on each of  $k$  random variables. This is often called a *data matrix*. What does the matrix  $X'X$  contain?
  - (ii) Show that the sample covariance matrix can be cast in matrix-vector form as:

$$\Sigma = \frac{1}{N} X'X - \bar{x}\bar{x}'$$

where  $\bar{x}$  is the vector containing the sample means.

5. Following the practicum of Section 9.5, you could try the following:
- (i) Pick a set of five stocks that you think will show nicely complementary properties for purposes of diversification. Then repeat the procedures of the text; compute the CAPM portfolio and see if the portfolio dominates that of the text, in the sense that it has a higher mean return and a lower variance.
  - (ii) Now get more ambitious and repeat the text exercise but with all 14 stocks on board. The hardest part—aside from darting back and forth across the spreadsheet—is to assemble the covariance matrix of returns. To speed things up, you might try the implied procedure in Q4 above, and do the whole assemblage in matrix -vector operation.

Has the extra diversification improved the expected return- risk trade off, for either of the portfolios—the Minimum Variance or CAPM?





# *Part Five*

## APPENDIX

# General Mathematical Review

## Appendix

### A1 Order of operations

As it is frequently necessary for the financial manager to perform chains of calculations, it is important to know the correct order in which to execute the various types of operations.

For example, if you are required to evaluate:

$$A = 8 - 4 \times 3 \div 6 + 2 \times 2^2,$$

where do you begin? Note that this is not just a simple matter of working from left to right! There are standard rules governing the order in which we perform the various arithmetic operations.

#### Rules

1. Calculations within parentheses are performed first.
2. Powers operations come next.
3. Divide and multiply are next.
4. Then add and subtract operations follow.

Where no brackets appear in the original equation, it is often easiest to start by grouping the  $\div$ ,  $\times$  and powers operations using parentheses.

#### Example A.1

$$\begin{aligned} A &= 8 - 4 \times 3 \div 6 + 2 \times 2^2 \\ &= 8 - (4 \times 3 \div 6) + (2 \times 2^2) \\ &= 8 - (2) + (2 \times 4) \\ &= 6 + 8 \\ &= 14 \end{aligned}$$

**Example A.2**

$$\begin{aligned}
 B &= 1 \div 3 \times (4 + 5)^2 - 7 \\
 &= 1 \div 3 \times (9)^2 - 7 \\
 &= (1 \div 3 \times 81) - 7 \\
 &= (27) - 7 \\
 &= 20
 \end{aligned}$$

**A2 Multiplication and division with signed numbers**

Numbers in calculations may be either positive or negative. When performing multiplication or division, if the two operands have the same sign the answer is positive. If they have different signs, then the result will be negative.

positive  $\times$  positive = positive  
 positive  $\div$  positive = positive  
 negative  $\times$  negative = positive  
 negative  $\div$  negative = positive  
 negative  $\times$  positive = negative  
 negative  $\div$  positive = negative

**Example A.3**


---

Find C for:

$$C = -16 \times -4 - 2 \times -6 + 3 \times -2 \times -5$$


---

First group the operations of multiplication and division using parentheses:

$$C = (-16 \times -4) - (2 \times -6) + (3 \times -2 \times -5)$$

Solve for sign within each set of parentheses:

$$C = (16 \times 4) - -(2 \times 6) + (3 \times 2 \times 5)$$

Perform the calculations within each set of parentheses working from left to right:

$$\begin{aligned} C &= 64 + 12 + 30 \\ &= 106 \end{aligned}$$

#### Example A.4

Calculate:

$$D = \frac{-24 \div -2 \times 4 - 5 \times -8 \div -2}{-12 - (5 \times -1)}$$

Group multiplication and division:

$$D = \frac{(-24 \div -2 \times 4) - (5 \times -8 \div -2)}{-12 - (-5)}$$

Solve for sign within the brackets:

$$D = \frac{(24 \div 2 \times 4) - (5 \times 8 \div 2)}{-12 + 5}$$

Perform the calculations within each bracket working from left to right:

$$\begin{aligned} D &= \frac{(48) - (20)}{-7} \\ &= \frac{28}{-7} \\ &= -4 \end{aligned}$$



## Exercises

**Find:**  $E = \frac{14 \div (-7) + 6 - 3 \times 4}{12 \div 6 - 4}$

$$F = \frac{4 \times 6 - (-12 \div (-2)) + 2 \times 3}{12 \div 3 + 5}$$

## A3 Powers and indices

*Powers, indices* and *exponents* are equivalent terms describing a shorthand way of writing repeated multiplication. They are frequently used in compound interest calculations.

### Example A.5

$$2 \times 2 \times 2 \times 2 \times 2 = 2^5$$

Read as “2 to the power 5” where 5 is called the *index, power* or *exponent* and 2 is the *base*. In the above example the base “2” is multiplied by itself five times.

## Rules for indices

1. Numbers with indices can only be added or subtracted when they have the same *base* and the same *index*.

### Example A.6

$$a^3 + a^3 = 2a^3$$

### Example A.7

$$b + 6b + 8b = 15b$$

**Example A.8**

$$\begin{aligned}
 E &= 8x^2 + 7y^2 + 3x^3 - 4y^2 - 2x^2 \\
 &= (8x^2 - 2x^2) + (7y^2 - 4y^2) + 3x^3 \\
 &= 6x^2 + 3y^2 + 3x^3
 \end{aligned}$$

2. Numbers with the same base but different powers can be multiplied. The result is the sum of the powers.

**Example A.9**

$$\begin{aligned}
 a^3 \times a^2 &= (a \times a \times a) \times (a \times a) \\
 &= (a \times a \times a \times a \times a) \\
 &= a^5
 \end{aligned}$$

that is:  $a^3 \times a^2 = a^{3+2} = a^5$

3. Division with Power Functions (Indices)

**Example A.10**

$$F = a^4 \div a^3$$

$$F = \frac{a \times a \times a \times a}{a \times a \times a}$$

When all the numbers in the numerator and denominator are multiplied, it is possible to reduce the equation as follows:

$$F = \frac{\cancel{a} \times \cancel{a} \times \cancel{a} \times a}{\cancel{a} \times \cancel{a} \times \cancel{a}}$$

(Note that this does not work if there is addition or subtraction in the numerator or the denominator.)

In the above example, we are left with the simplified form:

$$F = a$$

That is:

$$a^4 \div a^3 = a^{4-3} = a^1 = a$$

**Rule:** When dividing numbers with the same base but different powers, subtract the indices.

#### 4. Powers of Powers

When a number is raised first to one power (index), and then to another power, the indices are multiplied.

##### Example A.11

$$G = (a^3)^2 = (a^3) \times (a^3) = a \times a \times a \times a \times a \times a$$

$$\therefore (a^3)^2 = a^{3 \times 2} .$$

##### Example A.12

$$H = (a^2)^3 = a^2 \times a^2 \times a^2 = a^{2+2+2} = a^{3 \times 2} .$$

##### Example A.13

Caution:

$$a^{(2^3)} = a^8 \neq (a^2)^3$$

## 5. Negative Powers

A **non-zero base** raised to a negative power is the reciprocal of the base raised to the same (positive) power.

**Example A.14**

$$10^{-1} = \frac{1}{10^1} \qquad 2^{-5} = \frac{1}{2^5} \qquad x^{-n} = \frac{1}{x^n}$$

## 6. Zero Power

Any **non-zero base** raised to the power of **0** (zero) has the value **1**. This is because:

$$1 = \frac{a^n}{a^n} = a^0 ; \text{ so } a^0 = 1 .$$

## 7. Fractional Indices

Problems with fractional indices are common when you are trying to compare interest rates quoted with different compounding periods to determine their effective return.

**Rule:** The value of a positive base raised to a fractional power,  $n$ , is equal to the  $n^{\text{th}}$  root of the base.

$$x^{1/n} = \sqrt[n]{x}$$

$$10^{1/2} = \sqrt{10}$$

**Example A.15**

Solve for  $x$ :

$$(1 + x)^4 - 1 = 0.16$$

Add one (1) to both sides:

$$(1 + x)^4 = 1.16$$

Take the 4<sup>th</sup> root of both sides:

$$1 + x = 1.16^{1/4} = 1.16^{.25} = 1.03780199$$

Subtract one (1) from both sides:

$$x = .03780199$$

Expressed as a percentage:

$$x = 3.780199\%$$

*Note on Simplifying Fractions*

To simplify expressions with a fraction in the denominator, use the “invert and multiply” rule.

$$\frac{3}{4} \div \frac{4}{5} = \frac{3}{4} \times \frac{5}{4}$$

Using the “invert and multiply” rule this can be rewritten as:

$$\frac{3}{4} \times \frac{5}{4} = \frac{3 \times 5}{4 \times 4} = \frac{15}{16}$$

## A4 Logarithms

The logarithm, with base  $b$ , of a positive number  $A$ , is the power  $y$  such that  $b^y = A$ . The logarithm is written as  $\log_b A$  so that  $y = \log_b A$  means  $b^y = A$ .

### Examples

$$\log_{10} 100 = ? \quad 10^2 = 100 \quad \text{therefore } \log_{10} 100 = 2$$

$$\log_2 32 = 5 \quad \log_5 25 = 2 \quad \log_3 27 = 3$$

$$\ln x = \log_e x$$

$\ln$  is the notation for natural logs. They have as a base, the important mathematical constant  $e$ , the value of which is approximately 2.71828. These days, natural logs are almost universally used for calculations employing logarithms.

### Rules

1. The logarithm of a product of two or more positive numbers is the sum of the logarithms of the numbers. For example:

$$\log (A \times B) = \log A + \log B$$

2. The logarithm of a quotient of two positive numbers is the logarithm of the numerator minus the logarithm of the denominator. For example:

$$\log (A / B) = \log A - \log B$$

3. The logarithm of a positive number raised to a power is the product of the power and the logarithm of the number.

### Examples

$$\log (A^y) = y \log A$$

$$\log (x^2 / 3) = \log x^2 - \log 3 = 2 \log x - \log 3$$

Rule 3 is useful when solving for the power.

$$\text{Since } \log(x^y) = y \log x, \text{ then } y = \frac{\log(x^y)}{\log x}$$

Logarithms are extremely useful for finding answers to questions starting, “How long...?”

### Example A.16

How long does it take to double your money at an interest rate of 12% p.a. compound?

We wish to find  $x$  in the equation:

$$\begin{aligned} 1 \times (1 + .12)^x &= 2 \\ (1 + .12)^x &= 2 \end{aligned}$$

Take natural logs:

$$\begin{aligned} x \ln 1.12 &= \ln 2 \text{ (applying rule 3)} \\ x &= \frac{\ln 2}{\ln 1.12} = 6.116 \text{ years} \end{aligned}$$


## A5 Calculus

### Derivatives and differentials the easy way

Figure 1 below is a sketch of some more or less arbitrary function  $f(x)$  of a variable  $x$ . The value  $y = f(x)$  appears on the vertical axis. We are going to look at the behavior of the function  $f$  in a neighborhood of a point indicated as  $x$  on the sketch. Specifically, consider a small perturbation

$$\delta x = x_1 - x$$

At the point E, a tangent to the function has been drawn; this is the line ET. The point E has the vertical coordinate  $y = f(x)$ . The point Y has vertical coordinate  $y_1 = f(x_1)$ .



**Figure A.1 Differential coefficient as the slope**

The slope of the function  $f(x)$  at the point  $x$  is a magnitude denoted – in differential calculus – by the symbol:

$$\text{slope} = f'(x)$$

There are techniques for working out the slope  $t$  any given point  $x$ . We will give an example below. From basic geometry, the distance  $AT$  is equal to

$$AT = f'(x) \cdot \delta x$$

Notice that  $AT$  is almost, but not quite, equal to the distance  $AY$ . The latter is equal to  $y_1 - y = \delta y$ , where  $y_1 = f(x_1)$ . Thus we have:

$$\delta y \approx f'(x) \delta x,$$

where the wavy equals sign means “approximately equal to.”

You can see that the smaller the intervals  $\delta x$ ,  $\delta y$ , the closer the approximation will be. Indeed, the slope is defined by the limiting value of the ratio  $\delta y / \delta x$  as the distance  $\delta x$  tends to zero. Of course as this happens,  $\delta y$  will also tend to zero, but the *ratio* of the two will not in general tend to zero, and will become equal to the slope at the point  $x$ .



Conventionally, we use a special notation for the quantities  $\delta x$ ,  $\delta y$  as  $\delta x \rightarrow 0$ . We replace  $\delta x$  by a thing called  $dx$ , and  $\delta y$  by something called  $dy$ , and we refer to them as *differentials*. You might ask, “Why bother? Both are technically zero,” and this would be true. They are of the form of convenient fictional devices. If we know the slope  $f'(x)$  at a particular point  $x$ , then we can use the equality:

$$dy = f'(x) dx, \quad (1)$$

with the understanding that this represents the limit of an approximating process, in the course of which the ratio of  $dy$  to  $dx$  becomes equal to the slope of the function:

$$\frac{dy}{dx} = f'(x) \quad . \quad (2)$$

Differentials are heavily used in the mathematics of Ito processes (Chapter 8), again as a convenient shorthand. Indeed they are used in situations where the slope does not technically exist! In all cases, just think of them as small changes.

Techniques used to find the slope form the subject matter of *differential calculus*. There are certain basic results from which additional rules can be derived. For example:

$$\text{If } f(x) = a x^n \text{ then} \quad (3a)$$

$$f'(x) = a \times n x^{n-1} \quad (3b)$$

To differentiate a power of  $x$ , you just lower the degree by one (1), and put the original power out front. Constants like  $a$  in the above example make no difference.

### Differential equations

Sometimes  $x$  can stand for time, so that we are dealing with functions of time. In this case we would write  $f(t)$ . Suppose that someone said to you, “I know that the slope of a function, as  $t$  varies, is given by:

$$f'(t) = 12 t^2 \quad (4a)$$

Can you tell me what the original function must have been?”

From expressions (3a), (3b), and employing  $t$  in place of  $x$ , the answer is:

$$f(t) = 4t^3 + c \quad (4b)$$

How did we get this, and what does the constant  $c$  mean, anyway? The first term on the right hand side of (4b) was obtained by reversing (3b) to get the equivalent of (3a). In other words, if you differentiate the term  $4t^3$ , you get  $12t^2$  which is the given expression (4a). The process of reversing the differentiation is called taking the *anti derivative*, or more loosely, *integration*.

So far, so good. But how did the constant “ $c$ ” sneak into the answer, and how do we cope with it? When you differentiate a constant like  $c$ , the answer is always zero (ask yourself: what is the slope of a constant?). We could have started with any constant  $c$  in (4b) and still ended up with the given (4a). So we will have to allow for the possibility that there might have been some constant there anyway. To cope with the unpleasant feature of having some sort of indeterminate constant hanging around in our solution, we use any other information that might be available about the proposed function.

Suppose an informant had also told us that he knew that at time  $t = 0$ , the value of the function was  $f(0) = 10$ . For example, what was given to us as expression (4a) could be some sort of law of growth, referring to the rate of change of weight of Godzilla. At birth, which is  $t = 0$ , it is known that Godzilla weighed ten tons. Once we know this information, we can set  $t = 0$  in our proposed solution (4b) which now tells us that the only qualifying constant is  $c = 10$ . So the complete solution for the weight of Godzilla at any time  $t$  is:

$$f(t) = 4t^3 + 10$$

We have solved the *differential equation* (4a). To do so, we have used as a piece of further information that the value of  $f(t)$  at  $t = 0$  is 10. The latter is a *boundary condition*, or in this case, a particular form of boundary condition called an *initial condition*.

### **Integration** (briefly)

The antiderivative of a function is a much more general notion than the above treatment suggests. If you take an arbitrary function  $f(x)$ , and antidifferentiate it, you end up with the integrated function  $F(x)$ . This has the interpretation that at any

point  $x$ ,  $F(x)$  represents the area under  $f(x)$  to the left of the point  $x$ . Thus, integration is associated with areas.

You can see where this might have application. If we have a formula for a probability density and integrate, or antidifferentiate it, you end up with the corresponding distribution function, which is the cumulative area underneath the density.

### Partial derivatives

Sometimes you may have a function of two variables rather than one. For example, for set values of parameters such as the strike price, the interest rate, and the volatility, the value  $p$  of an option is a function of both the current stock price ( $S$ ) and the time to maturity ( $t$ ):

$$p = \pi(S, t)$$

We might be interested in the marginal effect of varying  $S$  while holding  $t$  constant. This is effectively the slope with respect to  $S$ . We might write it as:

$$\text{partial slope wrt } S = \frac{\partial \pi}{\partial S} .$$

This is called the partial derivative with respect to  $S$ . To find it, you simply pretend you are finding the ordinary derivative with respect to  $S$ , and treat  $t$  as a constant for such purposes. Similarly, we can write:

$$\text{partial derivative wrt } t = \frac{\partial \pi}{\partial t} .$$

If you are familiar with economics, you will be used to hearing that the total effect on the price  $p$  is a sum of the marginal effects, holding everything else constant. Suppose you observe small changes  $dS$ ,  $dt$  in the values of  $S$  and  $t$ . Then the total change in  $p$  is given by:

$$dp = \frac{\partial \pi}{\partial S} dS + \frac{\partial \pi}{\partial t} dt . \quad (5)$$

Once you have worked out the partial derivatives, you know what the total effect is by summing the partial effects. As defined in (5),  $dp$  is called the *total derivative* of the function  $\pi(S, t)$ .

Suppose a particular relationship between the partial derivatives were specified, for example, that the partial derivative with respect to  $t$  was equal to four times the partial derivative with respect to  $S$ . What function  $\pi$  would have such a property? This is an example of a *partial differential equation*. It is rather similar to an ordinary differential equation where there is only one argument varying, but this time there are two. It is often possible to find a particular function  $p$  which would have the required property. However the boundary conditions assume even greater importance, for there may be many functions that have the required property, but only one that satisfies the designated boundary condition. For partial differential equations, the boundary conditions often take the form of values of the function for an arbitrary  $S$  at a designated specific value  $T$  of  $t$ , the terminal value. For instance, the boundary condition for a call option is: At maturity  $t = T$ ,

$$\pi = \max ( S (T) - X , 0 )$$

For  $t < T$ , the call option price  $p$  does satisfy a partial differential equation, though it is a bit more complicated than the simple one mentioned above. But many other derivative prices actually satisfy exactly the same equation, and it is necessary to specify the boundary condition in the form given, in order to isolate the option price.

This completes our tour of calculus and differential equations.

# Index

## A

Absolute addressing 131  
Annuity 32, 33, 36, 37, 49, 59, 70, 140, 141  
Annuity due 32, 33, 36, 38, 39  
Arbitrage 74, 79, 81, 83, 85, 88, 122, 217  
ARCH 204, 224, 225  
Ask 104, 107, 108, 111–113, 121

## B

Bankers' Acceptances 24  
Bid 111, 112, 121  
Bivariate data 177  
Bond equivalent yield 12, 14, 24, 25  
Bootstrapping 75

## C

CD 24, 25, 27, 55, 56, 71, 154–156, 163, 249  
Commercial paper 12, 25  
Commodity currency 104–113, 119  
Compounding periods 10  
Continuous time 44–46, 97, 215, 218, 222  
Coupon 8, 13, 32, 47, 49–53, 59, 65, 70, 77, 78, 87, 89, 91, 93, 99, 100, 128,  
130, 138–140, 142, 143, 146  
Covariance 244, 246–248, 250–252, 254, 256, 260  
Cross-rates 113

## D

Degrees of freedom 170  
Density function 152, 165  
Direct Quotes 106  
Distribution 152, 153, 156, 159, 161, 166–173, 204, 213, 217, 226, 244, 246, 248  
Distribution function 152, 164, 165, 167, 219, 244, 277  
Drift 204, 215, 216, 218, 228  
Duration 31, 32, 57–61, 70, 100, 101

## **E**

Effective annual interest rate 32, 40  
Effective interest rate 8, 9, 32, 42, 45, 47, 59  
Equivalent rates 42  
Eurodollars 25  
Exchange rate 103–105, 107–114, 118, 120, 122

## **F**

Fed Funds 26  
Fixed interest securities 47  
Forward 33, 61, 73, 75, 76, 80, 82, 84, 95, 96, 98, 100, 103, 104, 114,  
116–120, 122, 206, 217  
Forward rates 81, 84, 86, 94, 96–98, 104, 114, 116  
Future value 11  
Futures 2, 12, 31, 33, 61–69, 71, 82, 84, 89, 144, 206–210, 212  
    Eurodollar Time Deposit 63, 71  
    30-year T-bond 65, 66  
    U.S. T-Bills 24

## **G**

Generation process 152, 167, 170

## **H**

Hedge 64, 71, 74, 82, 100, 119, 120, 207–209, 212, 219, 229  
Histogram 152, 153, 156–159, 161, 164, 167, 213  
Holding period yield 15, 25, 27

## **I**

Implied forward 64, 73, 74, 79–81, 84, 86, 87, 97, 100, 118  
Implied volatility 204, 220–223  
Indirect quotes 106  
Inflation 8, 21, 22, 224  
Inversions 104, 108  
Invoice price 33, 52, 53, 62, 65, 66, 146  
IRR 33, 48, 74, 77, 87, 144  
Ito process 204, 215, 216, 218, 220

## **K**

Kurtosis 152, 159, 161–163, 172  
    Leptokurtic 152  
    Platykurtic 152, 161, 163

**L**

Linear interpolation 19, 28, 88  
Long 74, 81, 83

**M**

Macaulay duration 32, 57, 58, 70  
Martingale 204, 214, 220  
Matrix 155, 235–243, 245–249, 251–254, 256, 257, 259, 260  
Mean 97, 152, 153, 160–173, 204–207, 216–219, 223, 224, 228, 235, 243, 245–250,  
252, 254, 256, 258, 260  
Mid-rate 104, 107, 108, 117  
Moment 41, 152, 160, 161, 167  
Money market securities 10, 12, 14, 24

**N**

Nominal 9, 39, 247  
NOW accounts 18  
Null hypothesis 169  
Number of days to maturity 8  
Numeraire 104, 105

**O**

OLS 207

**P**

Population 152, 153, 160, 161, 164, 166–168, 170–173  
Present value 8, 12–14, 16, 17, 24, 25, 32, 33, 35, 36, 38, 39, 47–49, 57, 70, 77,  
118, 128, 135, 138, 140  
Present value of an annuity 33  
Present value of an annuity due 33, 36, 38  
Price 8, 12–15, 22, 24, 25, 27, 32, 34, 47–56, 59–65, 71, 74, 77, 78, 82–84,  
88, 97, 103–105, 107, 108, 116–121, 139–141, 144, 146, 167, 203,  
204, 206, 209, 210, 212, 215, 217–219, 221, 223, 224, 226–230, 277, 278

**R**

Random sample 152, 153, 168, 226  
Random walk 204, 214, 227, 229  
Real interest rates 20  
Regression 149, 208, 209, 211, 212, 223, 225  
Relative addressing 131  
Repos 12, 25  
Reverses 12, 25

## **S**

Sample 152, 153, 160–162, 164, 166–171, 173, 205, 214, 223, 226, 243, 248, 259, 260  
Short 74  
Sigma notation 51  
Simple interest 10, 11, 52  
Skewness 153, 159, 161  
Spot rate 74, 75, 79, 98, 116, 118, 119  
Spot rates 76, 80, 98, 104, 113  
Standard deviation 153, 160, 165, 168–170, 173, 208, 215–217, 228, 258  
Statistic 153  
Stochastic process 203–205  
Stochastic variable 204  
Swaps 87, 89  
Sweep accounts 18

## **T**

Term structure of interest rates 73, 77  
Terms currency 104–107, 109, 119

## **U**

U.S. T-Bill futures 61  
Unit-discount bond 74

## **V**

Variance 152, 153, 160, 163, 166–168, 170, 171, 173, 204, 208, 213–216, 220, 223, 224, 226–229, 235, 245–252, 254, 256, 258, 260  
Vectors 235, 236

## **Y**

Yield curve 19, 28, 71, 74, 77, 86, 89

## **Z**

Zero coupon bonds 54, 78  
Zero coupon rates 57, 73, 74, 76, 77, 79, 80, 86, 88, 99, 100