

Financial Decision-Making

Michael R. Roberts
The Wharton School, University of Pennsylvania

Current Draft Date: August 2, 2024

Contents

Preface	i
Acknowledgements	v
I Decisions Everyone Makes	vii
1 The Finance Framework	1
1.1 Value	1
1.2 Financial Value	2
1.3 Opportunity Cost	4
1.3.1 Intuition for Opportunity Cost Adjustment	7
1.4 Price vs. Value	10
1.4.1 Nonfinancial Value	11
1.5 The Essence of Finance	12
1.6 Key Ideas	13
2 Retirement Savings and the Value of Education	15
2.1 Cost of College	16
2.1.1 Cost of College in Today's Dollars	16
2.1.2 Cost of College in Tomorrow's Dollars	19
2.2 Saving for Retirement	19
2.2.1 Consumption in Retirement	21
2.2.2 The Nest Egg	22
2.2.3 The Savings Strategy	24
2.2.4 What if...?	25
2.3 Inflation	29
2.3.1 Real Rate of Return	29

2.3.2	Real Cash Flows	31
2.3.3	Nominal or Real?	33
2.4	Saving for Retirement with Inflation and Taxes	34
2.4.1	Consumption in Retirement	34
2.4.2	The Nest Egg	36
2.4.3	The Savings Strategy	41
2.4.4	Reactions and Comments	42
2.5	The Value of College	43
2.5.1	Skip College	44
2.5.2	Go to College	45
2.5.3	The Decision	46
2.5.4	The Costs and Benefits of Going to School	47
2.6	Problem Solving Tips	48
2.7	Key Ideas	49
2.8	Technical Appendix	51
2.8.1	Derivation of Shortcut Formulas - Growing Annuity and Perpetuity	51
2.9	Problems	52
3	Borrowing Money	65
3.1	Loans and Interest Rates	66
3.2	Compounding	69
3.2.1	Summary	72
3.3	College Payment Plans	73
3.3.1	Measuring Time in Periods	74
3.3.2	Measuring Time in Years	75
3.4	Paying for College with Student Loans	76
3.4.1	Borrowing to Go to School	77
3.4.2	The Cost of Borrowing Money to Go to School	78
3.5	Financing a Home	81
3.5.1	Mortgage Mechanics	81
3.5.2	Shopping for a Mortgage	85
3.5.3	Refinancing a Mortgage	89
3.5.4	Paying Down Your Mortgage Early	91
3.6	Should You Pay Cash for a Car?	96
3.6.1	The Correct Way to Approach the Problem	97
3.6.2	Where Did Keppel et al. Go Wrong?	98

CONTENTS

3.7	Leasing a Car	100
3.7.1	Lease vs. Buy and Borrow	101
3.7.2	Lease Payment Calculations	103
3.8	Credit Cards	104
3.8.1	Institutional Details	105
3.8.2	Ryan's Bad Idea	106
3.9	Key Ideas	108
3.10	Technical Appendix	108
3.10.1	Selecting a Mortgage Based on Time in the Home	108
3.10.2	Continuously Compounded Interest Rates	109
3.11	Problems	110
II	Decisions Most People Make	121
4	A Financial Perspective of Business	123
4.1	Investors	124
4.1.1	Shareholders	124
4.1.2	Debtholders	125
4.1.3	Fundamental Value Relation	126
4.2	Business Decisions	127
4.2.1	Fundamental Value Relation	127
4.3	Nonfinancial Considerations	128
4.4	Key Ideas	129
5	Project Viability	131
5.1	Decision Criteria	132
5.1.1	Net Present Value (NPV)	133
5.1.2	Internal Rate of Return (IRR)	134
5.1.3	Payback Period	136
5.1.4	Money Multiples	138
5.1.5	Summary	139
5.2	Free Cash Flow	140
5.2.1	Profits	141
5.2.2	Investments	143
5.2.3	Summary	146

5.3	Assessing a Stand Alone Project	146
5.3.1	Forecast Drivers	147
5.3.2	Incremental Earnings	148
5.3.3	Net Long-Term Investments	152
5.3.4	Net Working Capital Investment	154
5.3.5	Free Cash Flows and Valuation	156
5.3.6	Comments	157
5.4	Sensitivity Analysis	158
5.4.1	Market demand and pricing	159
5.4.2	Cost Assumptions	160
5.4.3	Key Value Drivers and Break Even Points	161
5.4.4	Targeting KPIs	164
5.4.5	Scenario Analysis	165
5.5	Key Ideas	165
5.6	Technical Appendix	166
5.7	Problems	168
6	Project Selection	185
6.1	Choosing Among Mutually Exclusive Projects	186
6.1.1	Comparing Cisco 1 and 2	186
6.1.2	Juniper 1 and 2	191
6.1.3	Huawei 1 and 2	192
6.1.4	Further Discussion	194
6.2	Choosing Among Projects when Facing Constraints	196
6.2.1	Cypress Technologies	197
6.2.2	Profitability Index	198
6.3	Choosing Among Projects with Different Lifetimes	202
6.3.1	Annuity Equivalent Cash Flow	203
6.3.2	When to Invest	204
6.4	Customer Lifetime Value (CLV)	207
6.4.1	The Basic Model	207
6.4.2	DCF in Disguise	209
6.5	Key Ideas	210
6.6	Technical Appendix	211
6.7	Problems	212

CONTENTS

7 Investing: Fixed Income	227
7.1 Savings Accounts and Certificates of Deposit	228
7.1.1 Savings Accounts	228
7.1.2 Certificates of Deposit (CDs)	230
7.2 What are Bonds?	232
7.2.1 Maturity or Term	233
7.2.2 Repayment Schedules	233
7.2.3 Security and Priority	234
7.2.4 Covenants	235
7.2.5 The Most Important Bonds? U.S. Treasury Securities	236
7.3 Investing in T-bills and STRIPS	239
7.3.1 Valuation	239
7.3.2 Prices and Yields	241
7.3.3 Yield Curve and the Term Structure of Interest Rates	244
7.3.4 Liability Driven Investment: Cash Flow Matching	247
7.3.5 Price Dynamics and Rate Risk	250
7.3.6 Holding Period Returns	253
7.3.7 Expected Returns vs. Realized Returns	257
7.3.8 Market Efficiency and Market Timing	258
7.4 Investing in T-Notes and T-Bonds	261
7.4.1 Valuation	261
7.4.2 Computing the Yield-to-Maturity	265
7.4.3 Price Dynamics	269
7.4.4 Holding Period Returns	271
7.5 Bond Arbitrage and Short Sales	275
7.6 Interest Rate Risk	279
7.6.1 Intuition	279
7.6.2 DV01	280
7.6.3 Duration	283
7.7 Liability Driven Investment and Immunization	285
7.7.1 Duration Matching	286
7.8 Default Risk	287
7.8.1 Credit Ratings	289
7.9 Inflation Risk	292
7.9.1 TIPS	293

7.10	Term Structure of Interest Rates	294
7.10.1	What Does the Term Structure Mean?	296
7.10.2	The Fed and Monetary Policy	297
7.11	Key Ideas	299
7.12	Technical Appendix	300
7.12.1	Forward Rates	300
7.12.2	Reconciling Interest Rate Risk Measures	302
7.12.3	Taxes on Zero Coupon Bonds	303
7.12.4	Taxes on Coupon Bonds	304
7.12.5	Proofs	305
7.13	Problems	306
8	Investing: Stocks	321
8.1	What is Stock?	322
8.1.1	Decision Making at the Company	322
8.1.2	Shareholder Cash Flows	323
8.1.3	Different Types of Stock	325
8.2	Buying and Selling Stock	327
8.2.1	Price Quotes	329
8.2.2	Order Types	330
8.3	Making and Losing Money - Returns	332
8.3.1	Annual Returns	333
8.3.2	The Impact of Dividends	335
8.3.3	The Impact of Stock Splits	336
8.4	Historical Data and Risk-Reward Estimates	338
8.4.1	Expected Returns	338
8.4.2	Volatility	339
8.4.3	Scaling Expected Returns and Volatility	341
8.5	Stock Valuation	342
8.5.1	The Starting Point	342
8.5.2	The Effect of Investment Horizon on Value	343
8.6	Valuation: The Dividend Discount Model (DDM)	345
8.6.1	Valuing Microsoft with the Gordon Growth Model	346
8.6.2	Valuing Exelon with the Gordon Growth Model	353
8.6.3	Learning from Prices	356
8.6.4	Value Decomposition: Assets in Place vs. Growth Options	358

CONTENTS

8.6.5	Price-to-Earnings (P/E) Ratio	359
8.6.6	The PEG Ratio: P/E and Growth	362
8.6.7	Valuing Microsoft with Multistage Dividend Growth Models	363
8.6.8	How Many Stages of Growth?	365
8.7	Total Payout Model	366
8.7.1	Valuing Microsoft (Again)	368
8.8	Flow to Equity Method	369
8.8.1	Terminal Value Growth and the Short-term Horizon	370
8.9	Key ideas	372
8.10	Technical Appendix	373
8.10.1	ROE and Retention Rate	373
8.10.2	PVGO and Profitable Investment	374
8.10.3	Taxes	374
8.11	Problems	375
9	Investing: Portfolios	391
9.1	Portfolio Returns	392
9.1.1	Microsoft and Ferrari Portfolio	392
9.1.2	Portfolio Weights	394
9.1.3	Short Positions	398
9.1.4	Common Portfolio Weights	399
9.1.5	Portfolio Rebalancing	400
9.1.6	Expected Returns	401
9.2	Portfolio Risk	403
9.2.1	Limits to Diversification	406
9.2.2	Two Types of Risk: Diversifiable and Non-Diversifiable	407
9.3	Risk-Return Tradeoff: Mean-Variance Frontier	409
9.3.1	The Microsoft and Walmart Frontier	409
9.3.2	Efficient and Inefficient Portfolios	411
9.3.3	Minimum Variance Portfolio (MVP)	413
9.3.4	Changing Correlation	414
9.4	A Frontier with Many Assets	416
9.4.1	A New Measure of Risk - Covariance	419
9.5	Adding a Risk-Free Asset to Our Portfolio	420
9.5.1	Temporary New Efficient Frontier	422

9.5.2	Sharpe Ratio	423
9.5.3	Efficient Frontier	425
9.5.4	Different Lending and Borrowing Rates	428
9.6	What Diversification Does <i>Not</i> Mean	431
9.6.1	Return Risk vs. Dollar Risk	431
9.6.2	Time Diversification	432
9.7	Key Ideas	434
9.8	Technical Appendix	436
9.8.1	2-Asset Portfolios	436
9.8.2	N-Asset Portfolios	437
9.8.3	Limits to Diversification	438
9.9	Problems	441
III Decisions Business People Make		455
10 Estimating the Cost of Capital - Factor Models		457
10.1	Investment Implications	458
10.2	Estimating Expected Returns	460
10.3	Estimating Apple's Equity Cost of Capital	462
10.3.1	Step 1: Estimate the Risk-Free Rate	462
10.3.2	Step 2: Estimate the Market Risk Premium	463
10.3.3	Step 3: Estimate Beta (β)	465
10.3.4	Step 4: Putting it all together	465
10.4	Context for Beta	466
10.4.1	Beta for the Market Portfolio = 1	466
10.4.2	Beta for the Risk-free Asset = 0	468
10.4.3	Systematic Risk versus Idiosyncratic Risk	469
10.4.4	Beta Variation	471
10.4.5	Debt Betas	473
10.4.6	Portfolio Betas	474
10.5	Bringing it all Together	475
10.5.1	Capital Market Line and Efficient Portfolios	475
10.5.2	Security Market Line	480
10.6	Risk-Adjusted Investment Performance	481
10.7	Other Factor Models	485

CONTENTS

10.7.1	Apple's Equity Cost of Capital Revisited	486
10.7.2	Parting Thoughts	488
10.8	Key Ideas	488
10.9	Technical Appendix	490
10.9.1	Adjusting Beta Estimates	490
10.10	Problems	490
11	Optimal Financial Policy	497
11.1	Market Value Balance Sheet	498
11.1.1	Enterprise Value and Excess Assets	500
11.1.2	Expected Returns and Beta	501
11.1.3	Summary	504
11.2	Financing a Home	506
11.2.1	Setup	506
11.2.2	The Capital Structure Choice	506
11.2.3	Future Cash Flows and Values	507
11.2.4	Returns	508
11.2.5	Summary	511
11.3	When is Financing Irrelevant for Value?	513
11.3.1	Perfect Capital Markets	515
11.4	Leveraged Recaps, Payout Policy, and Excess Assets	516
11.4.1	The Perfect Capital Markets Benchmark	516
11.4.2	Considering Market Imperfections	521
11.4.3	Summary	533
11.5	Project Financing's Impact on the Firm	534
11.5.1	Firm and Project Setup	534
11.5.2	The Perfect Capital Markets Benchmark	535
11.5.3	Considering Market Imperfections	541
11.6	Project Financing's Impact on the Project	547
11.6.1	Dell's MVBS and The Tablet Project	547
11.6.2	Target Capital Structure & WACC Valuation	548
11.6.3	Predetermined Capital Structure & Adjusted Present Value	550
11.6.4	Summary	552
11.7	Leasing	553
11.7.1	Different Flavors of Leases	554
11.7.2	A Commercial Truck Lease	555

CONTENTS

11.7.3	The Perfect Capital Markets Benchmark	556
11.7.4	Considering Market Imperfections	557
11.8	Key Ideas	561
11.9	Technical Appendix	562
11.9.1	Project WACCs	562
11.9.2	Cost of Capitals with Predetermined Debt Policy	562
11.9.3	The Effect of Default	563
11.9.4	Personal Taxes and Capital Structure	564
11.10	Problems	565
12	Corporate Valuation (Incomplete)	573
12.1	Business Model	574
12.2	Discounted Cash Flow (DCF) Analysis	576
12.2.1	Valuation Strategy	576
12.2.2	Short-Term Value	578
12.2.3	Long-Term Value	584
12.2.4	Total Value	585
12.2.5	Debt Policy and Equity Values	586
12.2.6	Financial Statement Analysis	589
12.3	Market Multiples and Comparable Companies	591
12.3.1	Valuing a Home	591
12.3.2	Conceptual Framework	594
12.3.3	Valuing VB Dental	599
12.4	Key Ideas	602
12.5	Technical Appendix	602
A	Personal Budgeting	603
B	Financial Accounting	621
C	Spreadsheet Tips	657

CONTENTS

Copyright

Copyright ©2020 Michael R. Roberts. All rights reserved. No part of this book may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author. For permission requests, contact the author at the address below.

Michael R. Roberts
The Wharton School
University of Pennsylvania
3620 Locust Walk, Suite 2400
Philadelphia, PA 19104
Email: mrrobert@wharton.upenn.edu
Phone: (215) 573-9780

Disclaimer This book is for educational purposes only and should not be construed as providing financial, tax, legal, or any other advice. References to third parties are for demonstrative or descriptive purposes only.

Preface

This book is for people who need to practice finance - a group consisting of everyone but infants and toddlers - and who have little to no knowledge of finance. We make financial decisions in our personal and professional lives all the time. These decisions can have profound implications. They impact our financial, emotional, and even physical well-being, as well as the well-being of our loved ones. While abundant, financial advice is rarely suited to our specific needs, often incomplete, and frequently incorrect. Finance is important. We should understand how to practice it and doing so does not take years on Wall Street or an advanced degree.

So, I wrote this book, which is based in large part on my lectures to undergraduates, MBAs, and executives that I have had the privilege of teaching at the University of California Berkeley, the Fuqua School of Business at Duke University, and, for the last 20 years, the Wharton School at the University of Pennsylvania. It's also benefited from numerous private engagements with alumni, executives, and firms in which we put into action the lessons in this book. Finally, my more recent work with high school students has also played a role in some of its content and organization. The book is meant to be engaging and informal but precise, emphasizing intuition and applications. Its goal is to empower people to practice finance in their personal and professional lives, and prepare them for further study as needed.

To do so, I wanted to avoid having people slog through a 1,000+ page encyclopedic tome. (I've often found textbooks are written more for the academics selecting them than the students using them.) I chose not to include some topics (e.g., real and financial options, futures) or dedicate entire chapters to others (e.g., international finance, mergers and acquisitions). Proofs of key results are left to appendices, and theoretical excursions are limited. For most, the cost of these exclusions is low. You don't need to see many proofs or much theory to be a skilled practitioner of finance much like you don't need to be an automotive engineer to be a great race car driver. For those wanting to learn more, there are many other resources with which you'll be well prepared to engage after reading this book.

Instead, I decided to emphasize real-world applications and the thought process for tack-

ling financial challenges that most of us face in our lives. This choice makes the book more challenging, but hopefully more valuable. To ease the cognitive burden, I decided to introduce key financial principles - cash flows and opportunity cost - at the start of the book with personal finance applications that are more relatable and relevant for a broader audience. I've found that students can more easily digest and engage with new financial concepts by introducing them in settings with which they are already familiar. The added benefit of this approach is highlighting just how costly poor financial decision making can be and how valuable the lessons of this text are at a personal as well as professional level.

I've also organized the book around a small set of financial principles that tie all the applications together. So, rather than presenting a collection of loosely connected topics, the book illustrates how the same principles are used in every financial decision and how financial lingo and jargon can hide this elegant simplicity. The emphasis on applying financial principles allows readers to see the connections between different applications and highlights that there is only "one" finance despite the many different applications - personal finance, corporate finance, investment management, etc. This emphasis also ensures that readers are able to adapt to an ever-changing financial environment in which new financial products and services are constantly appearing.

To further increase the book's usefulness, each chapter is accompanied by

- an Excel workbook containing all the computations and financial models discussed in that chapter,
- end-of-chapter exercises designed to reinforce concepts and introduce additional applications,
- slides for each chapter,
- business cases that mimic real life scenarios in which decision-makers find themselves (for instructors), and
- end-of-chapter exercise solutions and business case teaching notes (for instructors).

The most important and broadly relevant material is at the beginning of the book - Part I: Decisions Everyone Makes. *Everyone* high school age and up should know this material, which covers the basics of financial decision-making in personal settings - how to construct a retirement savings strategy; the value of a college education; how to finance a home, car, school, etc. The appendix of the book also includes a discussion of personal budgeting.

Part II: Decisions Most People Make examines financial decision-making in business settings and details different investments - fixed income, stocks, and portfolios of assets. Understanding financial decision making in a business setting - which is no different from that in a personal setting - is important regardless of our chosen profession. This knowledge enables us to participate in resource allocation decisions and communicate with business leaders focused on financial goals. Related, the products offered by financial institutions and securities issued by businesses and governments are investment opportunities for savers. Understanding the risks and rewards of these opportunities is critical for both businesses and consumers.

Part III: Decisions Business People Make is more relevant for current and aspiring financial professionals, as well as business leaders. It shows how the financial sausage is made - e.g., how to estimate the cost of capital, how financing impacts companies and their investors, and how to value entire companies for a variety of purposes including financial, planning and analysis (FP&A), mergers and acquisitions, and leveraged buyouts.

I've used the material in the book in a variety of ways for different courses, some of which are listed here.

- Chapters 1 through 12 are used for a 14-week, semester-long core Finance course taught to undergraduates and MBAs.
- Chapters 1 through 4 and 7 through 9 are used for a two-week, intensive program on personal finance.
- Chapters 1, 4, 5, 6, and 10 and the Accounting appendix are used for a one-week, intensive executive education program aimed at nonfinancial executives.
- Chapters 10, 11, and 12 are used for a one-week, intensive executive education program aimed at financial executives.

The book can be used in several ways as suggested by the following table. I use the entire book for a 14-week, one semester core Finance course here at Wharton. For a half 6-week, half-semester course on personal finance I've used chapters 1 through 3 and 7 through 9. For a one-week intensive executive education course aimed at nonfinancial executives I rely on chapters also used subsets of the chapters for shorter courses. For example, I've used chapters 1-4

The prerequisites for this books are minimal. Two appendices detail the necessary spreadsheet and financial accounting knowledge. The mathematical prerequisite for most of this

book is arithmetic. Basic algebra, probability, and statistics are handy for some parts of the book. Calculus and linear algebra can be found in some of the technical appendices. Mathematically, most everything in this book can be understood by a 12-year-old. I know this because I made my 12-year-old daughter read it. More important than mathematical prerequisites is an interest in learning finance, something I could not make my daughter have...yet. The best way to build that interest is to know that the knowledge and skills you acquire from this book will make you and your loved ones substantially better off.

Acknowledgements

This book was made possible because of the teachings, help, and support of many people beginning with my family: Andreea, Sophie, and Max. My parents, Leonard and Patty, brother, Matthew, grandparents, Lois and Donald Isaacs, and aunt and uncle, Letty and Jerry Roberts, shaped who I am. Finally, I have to recognize my dogs: French Fry, Gummy, Bandit, and Luna who were by my side at different times during the writing of this book.

My educational development was influenced by far too many to name. However, I would be remiss to not mention a small number who played a particularly instrumental role in my professional development. David A. Freedman and Tom Rothenberg, my advisors and friends, taught me how to think probabilistically, statistically, and economically. Peter DeMarzo, Hayne Leland, and Mark Rubinstein inspired my interest in finance and showed me how it could be taught well. Ken Singleton and Bruno Biais, my co-editors at the *Journal of Finance*, helped expand my knowledge through numerous discussions and debates over the many papers that came across our desks.

I have also learned much from my colleagues and co-authors including, but not limited to, Jonathan Berk, Jules van Binsbergen, Jacob Boudoukh, Michael Bradley, Alon Brav, Sudheer Chava, John Cochrane, Joao Gomes, John Graham, Cam Harvey, Urban Jermann, Christopher Knittel, Pete Kyle, Mark Leary, Mike Lemmon, Juhani Linnainmaa, Daniel McFadden, Andrew Metrick, Angela Merrill, Roni Michaely, Mitchell Petersen, James Powell, Manju Puri, Matt Richardson, Nick Roussanov, Mike Schwert, Rob Stambouli, Amir Sufi, Vish Viswanathan, Jessica Wachter, Toni Whited, Robert Whitelaw, Amir Yaron, Moto Yogo, Rebecca Zarutskie, Jaime Zender, and Jeff Zwiebel.

Detailed feedback from Leila Ashtaryeh and Ian Pedroza was particularly helpful in writing this book, as was the feedback from many students I was fortunate enough to meet during the last 20 years. As they learned from me, I learned from them.

Part I

Decisions Everyone Makes

Chapter 1

The Finance Framework

This chapter

- provides the intuition behind financial decision making,
- defines financial value,
- introduces opportunity cost,
- distinguishes between value and price, and
- distills finance into one equation comprised of nothing more than arithmetic operations.

If the discussion here seems elementary, good. Conceptually, finance is not difficult.

1.1 Value

Finance is about decision-making, which requires weighing costs and benefits. When the benefits are bigger than the costs, the decision is a good one; when smaller, a bad one. Let's formalize this idea.

$$\text{Value} = \text{Benefits} - \text{Costs}$$

This equation says that value is positive when benefits are bigger than costs, negative when they're smaller. Put this way, good decisions create value, bad decisions destroy value.

Consider whether or not to eat a chocolate bar. The benefits are that it tastes good. The costs are that it has a lot of sugar and isn't particularly healthy. When I think about having chocolate, I weigh these costs and benefits to determine how much, if any, I should

have. When the benefits are relatively large, eating chocolate creates value for me, where value is just some notion of enjoyment or pleasure. When the benefits are relatively small, eating chocolate destroys value for me, perhaps because of guilt or concern about my health.

For most small decisions, the process of weighing costs and benefits is automatic. Nonetheless, this example is illustrative. We make decisions by weighing the costs and benefits whose relative sizes determine the value of the decision to us and, ultimately, the action we take.

1.2 Financial Value

Financial decisions work the same way. Financial benefits consist of money we receive, **cash inflows**. Financial costs consist of money we pay, **cash outflows**. While simple, these ideas are so important that the silly pictures in figures 1.1 and 1.2 illustrating them are worth the space.



Figure 1.1: Benefits = Cash Inflows



Figure 1.2: Costs = Cash Outflows

Intuition suggests that financial value should be the difference between these cash flows,

$$\text{Financial value} = \text{Cash inflows} - \text{Cash outflows}.$$

Financial value is positive when we receive more money than we pay and negative when we receive less than we pay. There's a nice logic to this relation, though we'll see shortly that it's incomplete.

First, let's recognize *when* we receive and pay money by drawing a **timeline**, which is presented in Figure 1.3. A timeline shows cash inflows and outflows by time period,

which are indicated by the labels on top of the timeline and the subscripts on the cash flows. The periods could be a year, quarter, month, day, etc. If we subtract cash outflows from cash inflows, period by period, we get a net cash flow for the period or, more simply, *CashFlow*, which I'll often abbreviate with *CF* later in the book. A timeline is a simple and surprisingly useful visual aid for solving financial problems because it details both the timing and magnitude of cash flows. We will use this tool repeatedly throughout the text.




Figure 1.3: Timeline and Cash Flows

From the structure of our timeline, financial value can be expressed as follows.

$$Value_t = CashFlow_{t+1} + CashFlow_{t+2} + CashFlow_{t+3} + \dots \quad (1.1)$$

Equation 1.1 says that the value of a decision at time t is equal to the sum of all future cash flows, which are the money we receive less the money we pay. The ellipsis (“...”) at the end of the equation indicate that the cash flows could go on indefinitely, which is represented by the timeline’s arrow in figure 1.3.

For example, imagine a deal that will pay us \$100 a year for two years, starting next year. Equation 1.1 says that the value of this opportunity today at $t = 0$ is $\$100 + \$100 = \$200$. Put differently, a fair price for this deal according to equation 1.1 is \$200. Figure 1.4 illustrates a \$200 payment today followed by \$100 inflows over the next two years.



Figure 1.4: Timeline and Cash Flows

If this deal doesn’t smell quite right, then our intuition is spot on. The value relation in equation 1.1 is missing something.

1.3 Opportunity Cost

Giving someone \$200 today in exchange for receiving \$100 per year for the next two years is (typically) a bad deal. Why? If we had \$200 today, we could invest that money and have more than \$200 in the future. Parting with money today comes with a cost, called an **opportunity cost**, that reflects the money we forgo by not investing it. Put differently, making us wait to receive money is costly because we can't invest it today. And, the longer you make us wait, the costlier it is because we are forgoing more investment earnings.

Exactly how much we forgo depends on how we would invest the money. In other words, if we don't take the deal and instead invest our money in something else, how much would we earn? Table 1 presents some investments and their corresponding average annual returns. For example, investing \$100 in a Treasury bill, which is a short-term loan to the U.S. federal government, generates \$103.30 one year later, on average. Investing \$100 in small-cap stocks, i.e., small companies, will generate \$117.28 one year later, on average. We'll discuss these investments and returns in more detail later. For now, it suffices to understand why these investments have different returns.

Investment	Return
Treasury bills (30-day)	3.30%
Treasury notes (10-year)	5.11%
Corporate bonds (Baa-rated)	7.19%
S&P500	11.82%
Small-cap stocks	17.28%

Table 1: Average Annual Investment Returns: 1927 - 2021

A clue can be found in Figure 1.5. The figure shows what each investment would be worth had we invested \$100 in 1927 and left the money in that investment for the next 94 years.¹ For example, \$100 invested in small-cap stocks in 1925 would be worth over \$3.6 million in 2020, whereas that same \$100 invested in Treasury bills would only be worth \$2,083.

The reason Treasury bill investors are willing to accept lower returns is that their investment is less risky than other investments. Likewise, corporate bond investors earn less than stock investors because their investment is less risky. This risk can be seen in the vertical movements or jaggedness of each line. The blue line corresponding to the small-cap stocks'

¹The calculations assume that all distributions - capital gains, dividends, and interest income - are reinvested in the asset.

value is very jagged compared to the relatively smooth purple line corresponding to the Treasury bills' value. Small-cap stocks lost more than half of their value during the Great Depression in the late 1920s!

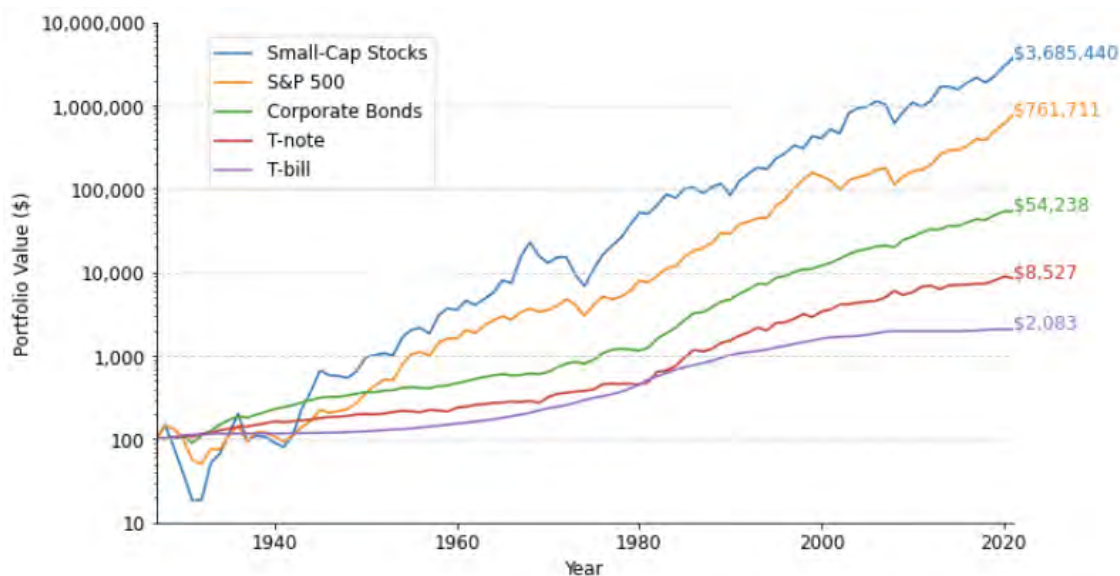


Figure 1.5: Value of \$100 Investments from 1927 to 2021

Bringing this discussion back to where we started, opportunity cost incorporates the *risk* of the opportunity. This is a fundamental point worth repeating. **The opportunity cost for a set of cash flows reflects the risk of those cash flows.** The riskier the cash flows, the higher the opportunity cost.

We might wonder if this means that cash flows that are guaranteed, i.e., risk-free, have an opportunity cost of zero. The answer is no. Also embedded in the opportunity cost is expected inflation or the expectation that the prices of goods and services will increase in the future. If the price of stuff goes up in the future, we'll need more money to buy that stuff. For example, if a pair of shoes costs \$100 today but \$110 one year from today, we'll need more money in the future to buy the shoes. So, the opportunity cost must account for expected price increases, as well as any risk associated with our investment.

Equation 1.2 shows how we account for opportunity cost in our expression of financial value.

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots \quad (1.2)$$

We've divided each cash flow by $(1+r)$ raised to the number of periods in the future in which it is received or paid.² The variable r measures our opportunity cost and is called

²Quick math refresher: raising quantities to a power is just shorthand for multiplication. For example,

a lot of things depending on the context: **cost of capital**, **discount rate**, and **expected return**. We'll see more names later.

Equation 1.2 says financial value at any time t is equal to the sum of future costs and benefits (i.e., cash flows) adjusted for their opportunity cost, r . When we compute the value today, $t = 0$, we often refer to financial value as **present value**, because it measures the value of an asset or decision as of the present moment.

Equation 1.2 is the key result we need to know to practice finance because it tells us what matters for value and how to estimate it.

Equation 1.2 shows that anything affecting future (i) cash flows or (ii) opportunity costs, r , will affect value. Likewise, value is simply a reflection of any and all future cash flows and their opportunity costs.

If we know the value of an opportunity, then we can make better decisions. For example, if the value of our mortgage has increased because interest rates have declined, then we should refinance it. If a stock is selling for \$20 per share, but is worth \$30 per share, then we should buy it. If a sales initiative requires a \$100 million investment but creates \$150 million of value, then we should invest in it. Most financial decisions boil down to an application of equation 1.2, which requires nothing more than basic arithmetic. What makes application of the value equation challenging is estimating the cash flows and opportunity cost, both of which occur in the future.

Let's reconsider our deal from the previous section by assuming that the risk of the cash flows is similar to investing in a corporate bond. From table 1, corporate bonds offer an average return of 7.19% per year. The present value of the deal (i.e., value today at $t = 0$) is

$$Value_0 = \frac{100}{(1 + 0.0719)} + \frac{100}{(1 + 0.0719)^2} = \$180.33.$$

We should pay no more than \$180.33 in exchange for receiving \$100 each year for the next two years when the opportunity cost of our money is 7.19%. If instead the risk of the cash flows was more accurately represented by an investment in small cap stocks, then the present value of the deal is

$$Value_0 = \frac{100}{(1 + 0.1728)} + \frac{100}{(1 + 0.1728)^2} = \$157.97.$$

$2^4 = 2 \times 2 \times 2 \times 2 = 16$. So, $(1 + r)^3 = (1 + r) \times (1 + r) \times (1 + r)$.

The deal is worth less to us today because the future cash flows are more uncertain. So, as the opportunity cost increases, indicative of riskier cash flows in the future, the value of the cash flows today decreases.

Because of its importance and centrality to everything we do, we'll refer to equation 1.2 throughout the book as our **fundamental value relation** and repeat it at the start of every chapter.

1.3.1 Intuition for Opportunity Cost Adjustment

For most, the rationale for the opportunity cost adjustment in our fundamental value relation is not immediately obvious. Why do we divide the cash flows by $(1 + r)^t$? Let's understand why with two analogies. First, imagine we have two boxes. One box weighs 125 pounds (lbs) and the other 59 kilograms (kg). How much do the two boxes weigh in total? What we *can't* do to answer this question is add 125 to 59 because they have different units. Saying the boxes weigh $125 + 59 = 184$ makes no sense; 184 what?

We have to convert the weights of the boxes into a common unit before adding them. It can be any unit - pounds, kilograms, ounces, tons, etc. - but it has to be the *same* unit. We can accomplish this by multiplying the weights by a conversion factor. Here are some examples.

$$\text{Weight in kilograms (kg): } 125 \text{ lbs} \times \frac{1 \text{ kg}}{2.205 \text{ lbs}} + 59 \text{ kg} = 115.69 \text{ kg}$$

$$\text{Weight in pounds (lbs): } 125 \text{ lbs} + 59 \text{ kg} \times \frac{2.205 \text{ lbs}}{1 \text{ kg}} = 255.10 \text{ lbs}$$

$$\text{Weight in ounces (oz): } 125 \text{ lbs} \times \frac{16 \text{ oz}}{1 \text{ lb}} + 59 \text{ kg} \times \frac{35.2 \text{ oz}}{1 \text{ kg}} = 4,076.80 \text{ oz}$$

Which conversion we perform depends on whether we're interested in knowing the total weight in kilograms, pounds, or ounces. Regardless, **to add (or subtract) numbers in a meaningful way, they must have the same units.** Additionally, we don't care if we have to carry boxes that weight 115.69 kg, 255.10 lbs, or 4,076.80 oz. They all weigh the same.

Now imagine we have 100 US dollars (\$100) and 100 Euros (€100). How much money do we have in total? Again, we have to convert the monies to a common currency, such as US Dollars, Euros, Pound sterling, etc. Here are some examples using conversion factors called

exchange rates from mid-December 2021.

$$\text{US dollars:} \quad \$100 + \text{€}100 \times \frac{\$1}{\text{€}0.88} = \$213.64$$

$$\text{Euros:} \quad \$100 \times \frac{\text{€}0.88}{\$1} + \text{€}100 = \text{€}188$$

$$\text{Pound sterling (£):} \quad \$100 \times \frac{\text{£}0.73}{\$1} + \text{€}100 \times \frac{\text{£}0.83}{\text{€}1} = \text{£}156$$

If you are unfamiliar with foreign currency, that's alright. As with adding weights, we can only add money if it is in the same currency units. Additionally, \$213.64, €188.00, and £156.00 are all the same value as of mid-December 2021, just expressed in different currencies.

In addition to a currency unit, money has a **time unit** indicating when we receive or pay it. Just like we can't add different currencies before converting them to a common currency, we can't add cash flows at different points in time until we convert them to a common time unit. Conversion to a common time unit is exactly what dividing by $(1 + r)^t$ does in our fundamental value relation (equation 1.2). The terms $1/(1 + r)^{t+s}$ for $s = 1, 2, 3, \dots$ act as exchange rates for time, converting the time unit of each future cash flow into period t time unit.

Moving Money Back in Time (Discounting)

Figure 1.6 shows the timeline for the deal discussed above. The first money we receive one year from now has a time unit of 1 and can be converted to time unit 0 by dividing by $(1 + 0.0719)^1 = 1.0719$. The \$93.29 in figure 1.6 is the present value of \$100 one year from now at an opportunity cost of 7.19%. Put differently, we are indifferent (i.e., equally happy) between receiving (or paying) \$93.29 today or \$100 one year from today because we can invest \$93.29 for one year at our opportunity cost of 7.19% and have \$100 at the end of the year, $\$93.29 \times (1 + 0.0719) = \100 . Using our box analogy, we don't care if we have to carry a box weighing 100 pounds or 45.25 kilograms, the weight is the same.

The second payment we receive two years from now is converted into time 0 units by dividing the cash flow by $(1 + 0.0719)^2 = 1.1490$. The value, \$87.03, is the present value of \$100 two years from today at an opportunity cost of 7.19%. Because the present values of each future cash flow have the same time units, period 0, we can add them to get the present value of the deal, \$180.32. (We're off by \$0.01 from the answer above because of rounding.) This process of moving cash flows *back* in time by dividing by $(1 + r)^t$ is called **discounting**.




Figure 1.6: Time Unit Conversion - Moving Money Back in Time

Moving Money Forward in Time (Compounding)

What if we want to know how much money we'll have three years from today if we invest our \$100 cash flows from the deal at 7.19%? Now we need to convert the time units on the cash flows to period 3 instead of period 0 (i.e., today). To move money forward in time we *multiply* the cash flows by $(1 + r)^t$ to get their **future values**. Figure 1.7 illustrates this process.




Figure 1.7: Time Unit Conversion - Moving Money Forward in Time

The figure shows how each cash flow's time unit is converted to a time unit 3, and the resulting future value. For example, \$100 one year from today is worth \$114.90 three years from today at an opportunity cost of 7.19%. Because the future values have the same time unit, period 3, we can add them.

$$\begin{aligned} Value_3 &= 100 \times (1 + 0.0719)^2 + 100 \times (1 + 0.0719) \\ &= 114.90 + 107.19 \\ &= \$222.09 \end{aligned}$$

The process of moving cash flows forward in time is called **compounding**.

Note, \$222.09 is the same number we get if we multiply the present value of the sum of the cash flows, \$180.33, by $(1 + 0.0719)^3$. This is not a coincidence. We can move sums of money around in time or each component of the sum. It doesn't matter as long as we always take care to only add cash flows with the same time unit.

Summary

We've introduced a few key concepts in this section worth summarizing before moving on.

1. Money has a time unit corresponding to when the money is paid or received. We can only add (or subtract) money if it has the same time unit.
2. The concept that the value of money depends on *when* it is paid or received is called the **time value of money**.
3. The quantity $1/(1 + r)^t$ is called a **discount factor** and it operates like a conversion factor or exchange rate for time. We use the discount factor to convert the time unit of money.
 - (a) Multiplying a cash flow by the discount factor (i.e., dividing by $(1 + r)^t$) moves money *back* in time by t periods. The process of moving money back in time is called **discounting**.
 - (b) Dividing a cash flow by the discount factor (i.e., multiplying by $(1 + r)^t$) moves money *forward* in time by t periods. The process of moving money forward in time is called **compounding**.

1.4 Price vs. Value

Price is what you pay for something. Value is what something is worth. They are not always the same.

For example, if we value a stock using our fundamental value relation (equation 1.2) at \$98 per share, and the stock is selling for \$98 per share then buying the stock doesn't make us any better or worse off. We would simply be paying the fair price for the stock. Yes, we may receive money in the future from dividends or by selling the stock at a higher price. But, the value of those future cash flows are accurately reflected in the current price. That's what our fundamental value relation does; it computes the value today of any future cash

flows. If instead the stock is selling for \$95 per share, then buying the stock today would make us better off. We would create $98 - 95 = \$3$ in value today for each share we purchase.

Likewise, if a company has an investment opportunity whose value is \$50 million but whose price or cost today is \$25 million, then undertaking the investment creates $50 - 25 = \$25$ million of value for the company today - a good idea. If the investment costs \$60 million today, then undertaking the investment destroys $50 - 60 = -\$10$ million of value for the company today - a bad idea. We can formalize this intuition by tweaking our fundamental value relation.

$$Net\ Value_t = -Price_t + \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots \quad (1.3)$$

We've made two changes to equation 1.2. We subtracted the current price ($Price_t$) from the future cash flows, and we changed $Value_t$ to $Net\ Value_t$ to indicate that the value is *after* accounting for the initial cost. Like present value, **net present value** or **NPV** refers to net value as of today ($t = 0$).

Let's again revisit our deal to receive \$100 a year for the next two years. We showed above that the value of the deal is \$180.33 when the opportunity cost is 7.19%. The net present value of the deal when the price is \$200 is $-200 + 180.33 = -\$19.67$, which is value-destructive for us. We wouldn't want to take this deal. Paying \$180.33 for the deal means the net present value is $-180.33 + 180.33 = \$0.00$. In this case, the deal neither creates nor destroys any value; it is fairly priced. However, if we could negotiate the price of this deal down to \$160, then our net present value would be $-160 + 180.33 = \$20.33$, thereby creating value and making us better off.

Equation 1.3 shows that we only *create* value (net value > 0) when we find opportunities in which the price is less than the present value of the future cash flows. When price is greater than the present value of future cash flows, we destroy value (net value < 0). (Actually, we'll see later how to create value by **short-selling** assets when the price is greater than the present value of future cash flows.) The key point is that only when price and value are different can we create or destroy value with our decisions. Otherwise, we're just getting a fair deal for the risk we're taking, and we're no better or worse off, financially speaking.

1.4.1 Nonfinancial Value

Value can be measured in many ways. Our example of eating a chocolate bar is one in which value is not measured in dollars and cents, ignoring the cost to buy the chocolate

bar. Rather, value is measured in feelings - happiness and anxiety - or health - blood sugar or cholesterol levels. Thus, it's best to think of financial value - the focus of this book - as complementary to other measures of value. Financial value is but one consideration in decision-making more broadly.

A great example of this complementarity is the decision to buy a house, which we'll explore in some detail later. From a finance perspective, buying a house is an investment decision with all sorts of interesting financial implications such as how to finance the purchase, its impact on our portfolio of investments, and the return from selling it. But, few people view the decision to purchase a primary residence solely through the lens of finance.

There can be great pleasure and pain associated with buying a house. These feelings matter. A house may make financial sense in that it will create financial value for us, but if we're going to be miserable living in the house, then that financial value is less important for our decision. Alternatively, if we find a house we love but can't afford, we may end up in financial straits and miserable. The point is that financial value is often only one consideration when making decisions. However, understanding the financial implications of our decisions is critical for keeping our feelings and emotions in check and avoiding pain and suffering later in life.

1.5 The Essence of Finance

Let's close the chapter by emphasizing the essence of finance contained in equation 1.2. Virtually every financial decision boils down to identifying the cash flows, the discount rate, or both. That's it. The primary concern of this book - of any book - on finance is in helping readers identify cash flows and discount rates in different contexts.

Depending on how one counts, the Wharton School at the University of Pennsylvania offers over 30 different finance courses that can appear completely unrelated. Some course titles include: Corporate Valuation, Venture Capital, Investment Management, Financial Derivatives, Real Estate Investments, Capital Markets, Distressed Investing, The Finance of Buyouts and Acquisitions, and ESG and Impact Investing

Indeed, if you were to take a financial engineer, someone specializing in financial derivatives, and place them in the middle of a corporate merger or a venture capital deal, they would almost surely be lost just as an investment banker or venture capitalist would likely be unable to assess a financial derivative. But the reason for their confusion has nothing to do with finance per se and everything to do with their inability to identify the relevant

cash flows and discount rates. The financial engineer, the investment banker, the venture capitalist - any practitioner of finance - are all trying to determine value, and value is always defined in the same way, equation 1.2. What's different is the context, which is critical for identifying cash flows and discount rates.

So, the real reason we, and many other schools, offer so many different finance courses is *not* to teach different versions of finance. There's only one finance! What these courses do is teach the different contexts in which finance is applied so that students can more easily identify the relevant cash flows and discount rates to estimate value.

This fact should be really comforting. The finance framework is simple. Equation 1.2 relies on nothing more than arithmetic. And, the framework is the same regardless of the financial question! What will challenge us is estimating the cash flows and discount rates in different applications. This is on what we'll spend most of our time.

1.6 Key Ideas

- Financial value, or present value when $t = 0$, is estimated with the following equation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

In words, financial value at any point in time, t , is equal to the sum of all the *future* financial costs and benefits (i.e., cash flows) adjusted for their opportunity cost (i.e., r).

- Net value, or net present value (NPV) when $t = 0$, subtracts the price or cost today from value. When net value is positive, we are creating value; when negative we are destroying value.

$$Net\ Value_t = -Price_t + \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

- The opportunity cost (a.k.a., discount rate, expected return, cost of capital), r , for a set of cash flows is the return on an otherwise similar risk investment. The greater the risk, the greater the opportunity cost.
- The expression “a dollar today is worth more than a dollar tomorrow” is a colloquialism for the time value of money and follows immediately from our fundamental value relation. Dividing cash flows by $(1+r)^t$ makes them smaller because we are dividing by a number larger than one (e.g., $1 + 0.05$). The further into the future, the bigger

the reduction in the value of the cash flow from the opportunity cost adjustment. Put differently, a dollar today is worth more than a dollar tomorrow and even more than a dollar two days from today.

- The discount factor, $1/(1+r)^t$, is our exchange rate for time. Discounting is the process of moving cash flows back in time t -periods by dividing them by $(1+r)^t$. Compounding is the process of moving cash flows forward in time t -periods by multiplying them $(1+r)^t$.
- Financial value is forward-looking, based entirely on what is *going to* happen, not what *has* happened. The challenge in practice is estimating the future cash flows and discount rates. Sometimes it's easy, often it's difficult. But, the difficulty should not mislead us into thinking that estimating financial value is a pointless exercise. The process of doing so is often more valuable than the end result because it forces us to clearly articulate why we are undertaking a decision.
- Financial value is one component of decision-making. Nonfinancial considerations such as feelings, health, etc., are relevant for the broader decision-making process. Our goal in this book is to ensure we are factoring financial considerations correctly in our decision-making.

Chapter 2

Retirement Savings and the Value of Education

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

This chapter

- reinforces our intuition for the Fundamental Value Relation through several applications,
- introduces inflation and taxes and shows how they affect our Fundamental Value Relation,
- presents several useful shortcuts for our Fundamental Value Relation when we are working with annuities and perpetuities, two cash flow streams that are common in practice, and
- applies our Fundamental Value Relation to answer several questions including:
 - How much money do we need to afford college?
 - How can we develop a retirement savings plan, and how do taxes and inflation affect that plan?
 - Does it make financial sense to go to college?

It's actually surprising - to me - that with nothing more than what we learned in the first few pages of this book, we can tackle these important questions. The one limitation in this chapter is the assumption that cash flows come and go on an annual basis. Fortunately, this isn't a terribly restrictive assumption. We can still answer important questions, and the only cost is some imprecision that has no effect on our decision making. We'll relax this assumption in the next chapter.

2.1 Cost of College

At the risk of redundancy, let's examine the cost of college in a manner similar to how we examined our friend's deal in the previous chapter. The goal is to reinforce the intuition for the Fundamental Value Relation and the time value of money in a realistic setting.

2.1.1 Cost of College in Today's Dollars

Tuition, fees, room, and board for undergraduates at the University of Pennsylvania total approximately \$80,000 per year in 2021 (holy schnike!). Assume this amount stays constant for all four years of school; unrealistic, but let's keep things simple. How much money do we need at the start of school in 2021 to cover all four payments, assuming the payments are made at the beginning of each year?

The first step to answering this, and most questions in finance, is to draw a timeline. Figure 2.1 shows our annual payments of \$80,000 per year for four years beginning at the start of each year. Period 0 corresponds to "today" or 2021 when we are trying to find value. We could have put negative signs in front of each cash flow to indicate cash outflows, but their meaning is unambiguous here.




Figure 2.1: School Savings Timeline

Ignoring opportunity costs, we would need $4 \times 80,000 = \$320,000$ at the start of school, period 0, to afford all four payments. However, because we can (and should!) invest any money we have, we'll need less than \$320,000. How much less depends on how we invest the money. Because we have to make the payments to complete school, any investment we take

Continuing with this logic, the third payment made two years from the start of school is converted into time 0 units by dividing the cash flow by $(1 + 0.05)^2 = 1.1025$. The value, \$72,562, is the present value of \$80,000 two years from today at an opportunity cost of 5%. Finally, \$69,107 is the present value of \$80,000 three years from today. Because the present values of each future cash flow all have the same time units, period 0, we can add them to get the cost of school when the opportunity cost of our money is 5%, \$297,859.84.

Another way to see what's happening in our value calculation is shown in table 1, called an **amortization table** or **amortization schedule**. **Amortization** refers to paying off debt in equal installments, as we do in paying off school, car, and home loans. Today, we put or **deposit** \$297,860 into a bank account earning 5% per year. We immediately take out or **withdraw** \$80,000 to make our first tuition payment. Thus, our bank account shows $297,859.84 - 80,000 = \$217,859.84$ today, the end of period 0.

The start of period 1 balance in our account is simply the previous end of period balance. For example, imagine we initially deposit our money at 12:00:00 pm on August 31, 2021, which corresponds to period 0. The start of period 1 would be immediately after the deposit is made, say one nanosecond later. In this very brief amount of time, we can safely assume that there is no time value of money effect. That is, the value of dollar is unchanged between nanoseconds. Hence, the start of period balance in the account is always equal to the previous period's ending balance. This is a convention we'll use in many applications in which we need to track how account balances - savings, loans, etc. - change over time.

During period 1, our savings earns 5% interest or $0.05 \times 217,859.84 = \$10,892.99$. We also must withdraw another \$80,000 for the second payment. Adding the interest to and subtracting the second withdrawal from the start of period balance leaves $217,859.84 + 10,892.99 - 80,000 = \$148,752.83$ in the account at the end of period 1. This process of earning interest and withdrawing tuition money continues for two more years, after which there is no more money left in the account.

Period	Start of period	Interest	Withdrawal	End of period
0				$297,859.84 - 80,000.00 = 217,859.84$
1	217,859.84	10,892.99	80,000.00	148,752.83
2	148,752.83	7,437.64	80,000.00	76,190.48
3	76,190.48	3,809.52	80,000.00	0.00

Table 1: Amortization Table for School Savings

2.1.2 Cost of College in Tomorrow's Dollars

If we chose not to go to college, how much money would we have in four years if we invested \$80,000 per year and earned 5%? To answer this question, we reverse the conversion process. Rather than *dividing* cash flows by $(1+r)^t$ to move future cash flows back in time, we *multiply* cash flows by $(1+r)^t$ to move cash flows forward in time. This process of multiplying cash flows by $(1+r)^t$ is called compounding, the result of which are future values. Figure 2.3 illustrates this process.




Figure 2.3: Time Unit Conversion - Moving Money Forward in Time

The figure shows how each cash flow's time unit is converted to a time unit 4 by compounding and the resulting future value. Because all of these future values have the same time unit, we can add them as in the following equation.

$$\begin{aligned}
 Value_4 &= 80,000 \times (1 + 0.05)^4 + 80,000 \times (1 + 0.05)^3 + 80,000 \times (1 + 0.05)^2 \\
 &\quad + 80,000 \times (1 + 0.05) \\
 &= 97,240.5 + 92,610 + 88,200 + 84,000 \\
 &= \$362,050.50
 \end{aligned}$$

Note, \$362,050.50 is the same number we get if we multiply the present value of the cash flows, \$297,859.84, by $(1 + 0.05)^4$.

2.2 Saving for Retirement

To illustrate just how much we can do with what we've learned so far, let's construct a retirement savings plan for someone we'll call Sophie. Her plan consists of answering three questions.

1. **Consumption in retirement.** How much money does Sophie need for expenses each year while in retirement when she's no longer working and earning money, and what is the expected return on her savings while in retirement?
2. **Nest egg.** How much money does Sophie need at the *start* of retirement to support her needs during retirement, a so-called "nest egg?"
3. **Savings strategy.** How much money does Sophie need to save each year while working to achieve her nest egg, and what is her expected return on her savings while working, if different from that in retirement?

We'll make some simplifying and admittedly unrealistic assumptions to begin. Doing so allows us to focus on the mechanics and get more comfortable with tools we've learned thus far. We'll relax some of these assumptions later and introduce some additional tools.

As always, we start with a timeline. To do so, we'll need some demographic information about Sophie. Assume Sophie just turned 30 years old and plans on retiring when she's 70. Based on actuarial life tables and her family history, Sophie expects to live to 87. Figure 2.4 shows the start of a timeline based on this information.




Figure 2.4: Retirement Savings Timeline

Period 0 as always corresponds to the decision date, today in this example when Sophie is 30 year old. She will work for 40 years and then retire when she turns 70. While working, she'll save money so that during retirement she can use those savings to take care of her needs - food, housing, utilities, travel, medical expenses, etc. A retirement savings plan consists of forecasts of her consumption needs in retirement, savings strategy while working, and investment strategy during both phases of life to get her opportunity cost. In other words, a retirement savings plan means we need to populate our timeline with cash flows, and identify the appropriate discount rate(s) for those cash flows based on how we plan to invest the cash flows.

Before doing so, let's recognize that we've already made some important assumptions. For example, we've assumed Sophie is going to work for 40 years. If she wants to work less

(or more), this will extend (shorten) her retirement years. Likewise, we've assumed she will live to 87, but what if she lives longer? Shorter? We'll explore these and other questions once we've got a baseline retirement plan in place. As we'll see, with a baseline model, it becomes almost trivial to answer these and many other "what if" questions, which amount to nothing more than changing a number in a calculation we've already performed.

2.2.1 Consumption in Retirement

How much money does Sophie need each year in retirement? To answer this, we need to know how much she'll need to spend on essentials, such as healthcare, food, housing, and utilities; and non-essentials, such as travel and entertainment. Considering this money is going to be spent 40 years from now, forecasting these figures is no small feat. However, this challenge does not mean the effort is worthless. On the contrary, the effort forces us to think very carefully about what we may need in retirement. While our forecasts are guaranteed to ultimately be wrong, we can always explore alternative forecasts - alternative scenarios - to understand the risks we'll face.

One starting point for this exercise is Sophie's current expenditures grossed up by **inflation**, or the rate of increase in the price of goods and services. Assume Sophie spends \$61,311.37 per year today on essentials and non-essentials. Historically, inflation has been approximately 3% per year in the U.S. When Sophie retires in 40 years, these same goods and services will cost $(1 + 0.03)^{40} - 1 = 226\%$ more than they do today, assuming future inflation is similar to the past. So, Sophie will need $61,311.37 \times (1 + 0.03)^{40} = \$200,000$ in 40 years to purchase the same goods and services she can purchase today for \$61,311.37.

One obvious shortcoming of this approach is that her consumption and expenditures will be very different 40 years from today. For example, medical expenses will become a larger part of her expenses in retirement. She may no longer have rent or mortgage payments. Her tastes and lifestyle will likely change. And so on. Additionally, future inflation could be quite different from historical inflation. That said, these possibilities are easily explored in our financial framework so we'll focus here on process and mechanics.

Let's assume Sophie needs \$200,000 during her first year of retirement. We'll also assume that Sophie needs \$200,000 *every year* in retirement to live the lifestyle she wants until she passes. In other words, we'll ignore the effect of inflation, which dictate that she needs more money each year so her consumption can keep pace with rising prices. This assumption is only to keep things simple for now. We'll relax this and other assumptions below. Figure 2.5 presents our timeline with these retirement cash flows.




Figure 2.5: Retirement Savings Timeline - Retirement Cash Flows

2.2.2 The Nest Egg

How much money does Sophie need at the start of her retirement, when she's 70, to ensure she will have enough money during her retirement years? If she stuffed the money under the bed, she would need $\$200,000 \times 17 = \3.4 million. Of course, this would be a bad idea for a couple of reasons. First, it's unsafe; someone could steal the money. Second, she would be missing out investment earnings. So, let's assume she plans on putting it in a savings account that earns 3% per year. She could invest the money in the stock market and expect to earn a higher *average* return, but because she has no income after retiring she plans on a conservative savings strategy that guarantees her a certain amount of money each year in retirement. Thus, the low investment return reflects the safety of the cash flows Sophie needs in retirement.

Armed with the cash flows, \$200,000 per year, and an expected return on her investments, 3%, we can use our Fundamental Value Relation (equation 1.2) to compute the value of her retirement needs at the start of her retirement in period 40 (age 70).

$$Value_{40} = \frac{200,000}{(1 + 0.03)} + \frac{200,000}{(1 + 0.03)^2} + \frac{200,000}{(1 + 0.03)^3} + \dots + \frac{200,000}{(1 + 0.03)^{17}} = \$2,633,223.69$$

Sophie needs approximately \$2.63 million at the start of her retirement. Figure 2.6 visualizes this calculation. Each cash flow in retirement is discounted to change the time unit to period 40 (age 70). Once they are all in the same time units, we can add them to get their value as of period 40, the start of Sophie's retirement.

Before continuing, we must emphasize an important aspect of financial analysis and modeling.

Never round numbers!

The cash flows in Figure 2.6, and other figures and tables throughout the book, are rounded only for presentation purposes. In the financial models and calculations, the numbers are never rounded. Small errors from rounding can quickly become large errors in calculations. Never round!




Figure 2.6: Retirement Savings Timeline - Nest egg

Annuities

Sophie's retirement needs - \$200,000 a year for 17 years - are what's known as an **annuity**, which is a stream of cash flows that are (i) constant, (ii) equally-spaced in time, and (iii) over a finite time horizon. There is a shortcut formula for computing the value of an annuity at any point in time, t , assuming the discount rate is also constant. (See the Technical Appendix for a proof of this result.)

$$\text{Value of an Annuity}_t = \frac{\text{CashFlow}}{r} \times (1 - (1 + r)^{-(T-t)}) \quad (2.1)$$

T is the number of cash flows, CashFlow is the constant cash flow, and r is the discount rate. Applying this result to Sophie's retirement needs yields the same amount we found above.

$$Value_{40} = \frac{200,000}{0.03} \times (1 - (1 + 0.03)^{-(57-40)}) = \$2,633,223.69.$$

Equation 2.1 is just a shorthand way of writing our Fundamental Value Relation (1.2) when the cash flows have the three features listed above - constant, equally-spaced in time, and a finite number of cash flows. This shortcut may seem unnecessary because we can easily compute and sum the present value of many cash flows in a spreadsheet. However, we'll see that the annuity relation in equation 2.1 can be used to answer many other questions not so easily answered by the brute force calculations in a spreadsheet.

When using equation (2.1) to find the value of an annuity, it is important to recognize that the first cash flow arrives one period in the future, as in Sophie's nest egg computation. We computed the value of Sophie's retirement cash flows as of period 40 (age 70), while her first withdrawal from the nest egg occurred one year later, in period 41 (age 71). In our school cost example (see Figure 2.1), this assumption is not met because the first payment of \$80,000 occurs immediately. However, we can still use the annuity result by treating

the future payments as an annuity and adding the present value of the annuity to today's payment. For our school cost example,

$$Value_0 = \underbrace{80,000}_{\text{Period 0 payment}} + \underbrace{\frac{80,000}{0.05} \times (1 - (1 + 0.05)^{-3})}_{\text{Present value of payments 1-3}} = \$297,859.84,$$

exactly what we computed above.

2.2.3 The Savings Strategy

How much does Sophie need to save each year while working to reach her target savings goal of \$2.633 million at the start of her retirement? Let's start by finding the present value of this amount by discounting from period 40 to period 0 (i.e., today) assuming Sophie will undertake an investment strategy while she works that is similar in terms of risk to that when she's retired. In other words, let's use the same discount rate we used for her retirement, 3%.

$$Value_0 = \frac{2,633,223.69}{(1 + 0.03)^{40}} = \$807,232.$$

Figure 2.7 illustrates how this calculation converts the \$2.633 million forty years from today into \$807,232 in today's dollars (i.e., present value).




Figure 2.7: Retirement Savings Timeline - Present Value of Nest egg

Another way to view the present value of Sophie's nest egg is as the amount of money she needs today to avoid having to save for retirement. If Sophie has \$807,232.73 today at age 30 and invests that money for the next 40 years at 3%, then her money will grow to be \$2,633,223.69 when she's 70 years old.

$$807,232.73(1 + 0.03)^{40} = \$2,633,223.69$$

No additional savings are needed.

1. What if Sophie had already saved \$250,000 by the time she turned 30? She could use this money to reduce how much she needs to save each year to reach her nest egg goal. How much? Because the \$250,000 has a time unit of 0, we can deduct this money from the present value of her nest egg, which also has a time unit of 0. Her new annual savings requirement is therefore

$$CF = \frac{(807,232 - 250,000) \times 0.03}{1 - (1 + 0.03)^{-40}} = \$24,107.21.$$

More precisely, to reach her nest egg target of \$2.63 million, Sophie needs to do two things. First, she needs to invest the \$250,000 she already has for 40 years at 3% per year. Second, she needs to save \$24,107.21 per year at 3% for 40 year.

2. What if Sophie wants to leave an inheritance (i.e., bequest) of \$500,000 for her kids when she dies? This inheritance is just another cash flow she needs in retirement in her last year of life, which is illustrated in figure 2.9 along with the \$200,000 annual cash flows fulfilling her consumption needs.




Figure 2.9: Retirement Savings with a Bequest

The nest egg needed to support Sophie and her kids' inheritance is

$$Value_{40} = \underbrace{\frac{200,000}{0.03} \times (1 - (1 + 0.03)^{-17})}_{Value_{40} \text{ of retirement needs}} + \underbrace{\frac{500,000}{(1 + 0.03)^{17}}}_{Value_{40} \text{ of bequest}} = \$2,935,731.92.$$

The first term in the equation computes the value as of period 40 of Sophie's retirement needs. The second term computes the value as of period 40 of the inheritance Sophie wants to leave for her children, i.e., her bequest.

The present value of this nest egg is

$$Value_0 = \frac{2,935,731.92}{(1 + 0.03)^{40}} = \$899,968.70.$$

To meet her new nest egg goal, Sophie needs to save

$$CF = \frac{899,968.70 \times 0.03}{1 - (1 + 0.03)^{-40}} = \$38,934.79$$

each year while working, assuming she has no other savings today and is able to earn 3% per year in interest.

3. What if Sophie wants to ensure that *all* of her descendants - kids, grandkids, great-grandkids,... - are able to withdraw \$200,000 every year *forever* from the estate she leaves them when she dies? The timeline with her retirement and post-death cash flows is illustrated in figure 2.10. There is no age indicated on the timeline after 87 because Sophie has passed on. However, the timeline continues forever because she wants all of her descendants to have \$200,000 every year.




Figure 2.10: Retirement Savings with a Perpetual Bequest

This question is not unlike the previous one in which Sophie leaves her kids a lump sum inheritance when she dies. The trick is figuring out the lump sum she should leave to ensure that everyone can withdraw \$200,000 a year forever.

Drawing down money forever might seem impossible, or it might seem to require an infinite amount of money. But, because cash flows are worth less and less the further into the future they arrive, Sophie can leave all her descendants a lump sum that is far from infinite, and still support them forever. Let's assume that the money she leaves everyone earns 3% per annum. When she passes at age 87, 57 years from today, she will have to leave her descendants a lump sum of money equal to the sum of the present values of all future cash flows.

$$\text{Value}_{57} = \frac{200,000}{1 + 0.03} + \frac{200,000}{(1 + 0.03)^2} + \frac{200,000}{(1 + 0.03)^3} + \dots$$

Clearly, it is impossible to compute this value by the brute force method of discounting each cash flow and summing because there are an infinite number of cash flows! Fortunately, math helps us out of this conundrum. A constant stream of cash flows, equally-spaced in time, and lasting forever is called a **perpetuity**. As long as the discount rate is greater than zero, this infinite sum is equal to the ratio of the cash flow to the discount rate, or

$$\text{Value}_t = \frac{\text{CashFlow}}{r} \quad (2.3)$$

Plugging the \$200,000 per year withdrawal and 3% discount rate into this equation produces

$$\text{Value}_{57} = \frac{200,000}{0.03} = \$6,666,666.67.$$

Leaving \$6.67 million invested forever at 3% when she passes ensures that Sophie's descendants can withdraw \$200,000 per year from the account *forever*.

4. *Challenging*: What if Sophie is concerned about outliving her money - living beyond 87 years? Let's say she decides to consume less in retirement, \$175,000 per year, and wants to know how much longer her nest egg will last? Remember, Sophie's consumption in retirement is just an annuity whose present value is given by equation 2.1. So, what's she's really asking is given her nest egg, Value_t , her consumption in retirement, CashFlow , and her investment earnings, r , how many years can she consume before running out of money, $T - t$?

Solving equation 2.1 for $(T - t)$, the length of the annuity, we get

$$T - t = -\frac{\ln\left(1 - \frac{\text{Value}_t \times r}{\text{CF}}\right)}{\ln(1 + r)} \quad (2.4)$$

where \ln is the natural logarithm.¹ Now, we can just plug and chug.

$$T - t = -\frac{\ln\left(1 - \frac{2,633,233.69 \times 0.03}{175,000}\right)}{\ln(1 + 0.03)} = 20.31 \text{ years}$$

By reducing her annual consumption in retirement to \$175,000 per year, Sophie can consume for 20.31 years, or a little over three years longer than if she consumes \$200,000 per year.

There are many other “what if” questions we could ask. What if Sophie decides to retire at 60 instead of 70? What if Sophie decides to invest more aggressively during her working years, say by investing in stocks that return on average 10% per annum? What if Sophie wants to save more each year as her income grows? With the tools we've learned thus far, we can easily explore how each change affects the size of her nest egg at the start of retirement and how much she has to save each year to reach that nest egg. Our next attempt at constructing Sophie's retirement plan explores some of these changes. But, first we'll expand on a key concept that was missing from our analysis.

¹For any real number x , $\ln(e^x) = x$, where e is the mathematical constant e , approximately equal to 2.71828.

2.3 Inflation

The end goal of financial decision-making for most isn't the accumulation of money per se. Money is just a means to an end, that end being the ability to purchase and consume goods and services - food, housing, vacations, etc. What ultimately matters for our welfare and happiness is our purchasing power, and more money typically means more purchasing power - e.g., better food, nicer housing, longer vacations, etc. We say typically because the growth in our money may not be able to keep pace with the growth in prices or **inflation**, which was introduced above. So, while investing money typically generates more money in the future, if inflation is higher than the return on our money then we have a problem. We won't be able to buy as much stuff in the future as we can today.

Consider oranges. In 2022, the average price per pound was \$1.70. Imagine economists expect the price of oranges to increase by 7% in 2023 to $1.70 \times (1 + 0.07) = \1.82 per pound because of increasing demand or decreasing supply. If we have \$1.70 today, we can buy a pound of oranges today. If we have a \$1.70 next year, we can't; we'll be short $1.82 - 1.70 = \$0.12$. What can we do? Well, we could (should) invest our money to earn a return.

Now imagine that we save \$1.70 this year and those savings earn 5%. The future value of our savings one year from today will be $1.70 \times (1 + 0.05) = \1.79 , still short of the \$1.82 we need to buy a pound of oranges. We've made money, but lost purchasing power. If instead orange prices were expected to increase by 3%, then the cost of oranges in 2023 would be $1.70 \times (1 + 0.03) = \1.75 per pound. Now we've gained purchasing power with our 5% investment return because we can not only afford the oranges, but we'll also have some money left over ($1.79 - 1.75 = \$0.04$). Notice that the relation between our investment return and inflation dictates whether or not we can afford oranges next year.

2.3.1 Real Rate of Return

The return on our money is called a **nominal return**. It measures how much money we earn. The return on our purchasing power is called a **real return**. It measures how much purchasing power we earn - how much more stuff we can buy. Let's denote the real return by rr , which is defined as follows.

$$rr = \frac{1 + r}{1 + \pi} - 1. \quad (2.5)$$

The nominal return is denoted by r and corresponds to the opportunity cost we've been talking about all along. Expected inflation is denoted by π (the Greek letter "pi"), which

corresponds to the expected increase in the prices of goods and services over the next year. It's important to clarify that when people talk about inflation, they are talking about increases in the *general price level*, an average price increase across many different goods and services. The price of any one good, like oranges, can increase (or decrease) slower or faster than average prices.

Figure 2.11 plots U.S. inflation from 1960 to 2022. There were several inflationary episodes during this period. Rising food and energy prices coupled with the end of the Nixon wage-price controls program led to the spike in 1974. Continuing high energy prices and real estate values lead to the spike in 1979. More recently, inflation increased following the Covid pandemic which led to broken supply chains and unprecedented fiscal stimulus. In 2008, inflation was negative, i.e., **deflation**, meaning the price level fell.




Figure 2.11: U.S. Inflation (Source: St. Louis FRED)

An often used approximation for equation 2.5 is

$$rr = r - \pi. \quad (2.6)$$

This approximation works well when nominal rates and inflation are small, say less than 10%.²

²This equation is based on a Taylor series approximation to equation 2.5.

Using the figures from our oranges example, the real return on our savings when our nominal return is 5% and the expected rate of inflation is 7% is

$$rr = \frac{1 + 0.05}{1 + 0.07} - 1 = -0.0187$$

While we are earning 5% on our money, prices are increasing by 7%, and our purchasing power is declining by 1.87%. We're making more money, but we can buy less stuff. If expected inflation is 3%, then our real return is

$$rr = \frac{1 + 0.05}{1 + 0.03} - 1 = 0.0194$$

or 1.94%. Now, we earn more than enough money to keep pace with increasing prices so we can buy even more stuff in the future than we could today.

The message here is really important, and will be repeated. A lower bound on any investment strategy should be to at least keep pace with inflation. Otherwise, we'll be able to purchase less and less stuff - food, housing, medical care, etc. - and our standard of living will deteriorate. In other words, we ultimately care about our real rate of return, not our nominal rate of return.

2.3.2 Real Cash Flows

Consider our cost of college example from above. Let's assume that the expected rate of inflation is 2% per year and the cost of college grows at the same rate. The timeline displaying **nominal cash flows** is shown in figure 2.12. Nominal cash flows correspond to what we will actually pay each year for schooling.

Period (years)	0	1	2	3	4
Nominal cash flows(\$000s)	80	$80(1 + 0.02)$ = 81.60	$80(1 + 0.02)^2$ = 83.23	$80(1 + 0.02)^3$ = 84.90	

Figure 2.12: School Costs with Inflation (Nominal Cash Flows)

Real cash flows measure the purchasing power of money relative to a base year. The term “real” as used here drives me nuts because real cash flows are anything but real. Unfortunately, we're stuck with this lingo. To compute the real cash flows relative to today, $t = 0$, we need to **deflate** the nominal cash flows by dividing them by $(1 + \pi)^t$, as shown in figure 2.13

Period (years)	0	1	2	3	4
Nominal cash flows(\$000s)	80	81.60	83.23	84.90	
Real cash flows(\$000s)	80	$\frac{81.57}{(1+0.02)}$ = 80	$\frac{83.17}{(1+0.02)^2}$ = 80	$\frac{84.80}{(1+0.02)^3}$ = 80	

Figure 2.13: School Costs with Inflation (Real Cash Flows)

The real cash flows measure the purchasing power relative to the base year. Because the cost of school is increasing at the same rate as inflation, the real cash flows are all the same and equal to today's cash flow, \$80,000. For example, the \$81,600 cost of college one year from today has the same purchasing power as \$80,000 today; it buys one year of schooling. More generally, imagine buying \$80,000 of "stuff" - goods and services - today. To buy that same stuff one year from today, we would need \$81,569. To buy that same stuff two years from today, we would need \$83,168. And, so on. So, the \$80,000 real cash flow each year is telling us the corresponding purchasing power in today's terms of the nominal cash flow each year (81,568, 83,168, 84,798).

Now assume that the cost of college increases by 4% per year. The timeline with nominal cash flows is shown in figure 2.14. These nominal cash flows tell us how much money we have to pay each year to go to Penn. The 4% increase in these nominal cash flows tell us that college is getting more expensive each year.

Period (years)	0	1	2	3	4
Nominal cash flows(\$000s)	80	$80(1 + 0.04)$ = 83.20	$80(1 + 0.04)^2$ = 86.53	$80(1 + 0.04)^3$ = 89.99	

Figure 2.14: School Costs Outpacing Inflation (Nominal Cash Flows)

Figure 2.15 shows the timeline with the corresponding real cash flows computed by **deflating** (dividing by $(1 + \pi)^t$) each nominal cash flow by a 2% rate of inflation. In other words, the cost of school is rising more rapidly (4% per year) than other goods and services in the economy. As a result, the real cash flows are all greater than \$80,000 and growing. In fact, this has been a long-running trend in higher education. Between 2002 and 2022, the cost of public college tuition has grown by 68% while the **consumer price index (CPI)** that measures the price level of all goods and services in the economy grew by only 39%.³

³Data on the consumer price index from the Bureau of Labor Statistics Consumer Price Index Inflation

Period (years)	0	1	2	3	4
Nominal cash flows(\$000s)	80	83.20	86.53	89.99	
Real cash flows(\$000s)	80	$\frac{83.20}{(1+0.02)}$ = 81.57	$\frac{86.53}{(1+0.02)^2}$ = 83.17	$\frac{89.99}{(1+0.02)^3}$ = 84.80	

Figure 2.15: School Costs Outpacing Inflation (Real Cash Flows)

High inflation is problematic. As mentioned above, if the prices of goods and services are increasing at a faster rate than our income or investment returns, then we can't buy as much stuff. Put simply, we're worse off over time even though we are making more money because that money can't buy as much stuff. This situation is exactly what occurred in 2021 and 2022 in the U.S. when inflation shot up to over 9% on an annual basis, while interest rates and wage growth were relatively low - approximately 4% per year for each.⁴ Many people couldn't earn enough money - working or investing - to keep up with increases in the price of goods and services.

2.3.3 Nominal or Real?

Should we use nominal or real returns and cash flows in our Fundamental Value Formula (equation 1.2)? It turns out, it doesn't matter. They will give us the same answer as long as we don't mix nominal and real measures.⁵ Specifically, **we discount nominal cash flows by nominal returns, and real cash flows by real returns. We never mix.**

Let's use our college example in which the costs increase by 4% per year and inflation is 2% to illustrate the equivalence. The nominal cash flows are shown in figure 2.14. With a nominal discount rate of 5%, the present value of the cost of college is

$$80,000 + \frac{83,200}{(1 + 0.05)} + \frac{86,528}{(1 + 0.05)^2} + \frac{89,989}{(1 + 0.05)^3} = \$315,457.53.$$

Calculator. Data on average public college tuition from Table 330.10, Average undergraduate tuition, fees, room, and board rates charged for full-time students in degree-granting postsecondary institutions, by level and control of institution: Selected academic years, 1963-64 through 2022-23, from the Digest of Education Statistics at the National Center for Education Statistics.

⁴Data sources include the Bureau of Labor Statistics, Federal Reserve Bank of Atlanta, and the US Census Bureau

⁵Mathematically, when using real values in the Fundamental Value Relation, we are dividing both numerator and denominator of each term by $(1 + \pi)^t$. These factors cancel to get nominal cash flows and returns.

The real cash flows for college costs are shown in figure 2.15. The real discount rate is $(1 + 0.05)/(1 + 0.02) - 1 = 0.0294\%$, implying that the present value of the cost of college is

$$80,000 + \frac{81,569}{(1 + 0.0294)} + \frac{83,168}{(1 + 0.0294)^2} + \frac{84,790}{(1 + 0.0294)^3} = \$315,457.53.$$

In practice, most people work with nominal cash flows and returns. An exception is in very high inflationary environments where nominal cash flows become extremely large even over short time horizons. Several countries have experienced episodes of **hyperinflation** or extraordinarily high inflation. In 1985, Bolivia experienced annual inflation of 20,000%. In the 2000s, Zimbabwe's economy was crushed by persistent hyperinflation that began at 624% in 2004 and reached 2,200,000% in 2008. The government printed 100 trillion Zimbabwe dollar bills. The Zimbabwe currency was all but worthless outside the country with an exchange rate of \$688 billion Zimbabwe dollars to \$1 US.

2.4 Saving for Retirement with Inflation and Taxes

Let's take another pass at Sophie's retirement savings strategy, this time incorporating more realism. Specifically, let's assume the following.

- Sophie needs \$200,000 at the end of her first year in retirement, after which her consumption needs to grow by 3% per year to keep up with expected inflation.
- She invests more aggressively during her working years, say by investing in stocks, and expects to earn 10% per annum before switching to a more conservative investment plan in retirement that earns 5%.
- She will grow her savings each year while working by 4% - her anticipated average annual salary increase.
- She will face a 30% tax rate on all income and a 20% tax rate on all investment earnings.

She still plans on retiring at age 70 and moving on at 87.

2.4.1 Consumption in Retirement

Figure 2.16 presents the timeline of Sophie's new retirement needs after accounting for inflation. Remember, these cash flows represent the money she needs each year in retirement

to purchase food, housing, clothes, vacations, medicine, etc. Because of inflation, the cost of these goods and services is expected to increase by 3% per year. So, the money Sophie needs in retirement will also have to increase by 3% per year. As a result, she will need over $200,000 \times (1 + 0.03)^{16} = \$320,941$ when she's 87.




Figure 2.16: Retirement Consumption with Inflation

Income Tax

Depending on the source from which her savings come, Sophie may be required to pay income taxes on her withdrawals. As of 2024, the tax rate varies across sources - Social Security, retirement accounts (e.g., **Individual Retirement Account/IRA** or **ROTH IRA**), and private savings - and varies across time with changes in tax policy. To avoid getting stuck in the weeds of tax law, we've assumed Sophie has to pay a constant 30% tax rate throughout her retirement. So, for each dollar she withdraws from her savings, she pays the government \$0.30.

For example, in her first year of retirement, Sophie needs \$200,000 for her anticipated expenses. Therefore, she needs to withdraw $200,000 / (1 - 0.30) = \$285,714$. When she does, she can pay $30\% \times 285,714 = \$85,714$ to the government and keep the remaining $285,714 - 85,714 = \$200,000$ to spend on expenses. More generally, each year the amount she needs to withdraw is given by the following expression.

$$\text{Pre-tax withdrawal} = \frac{\text{Post-tax money}}{(1 - \text{TaxRate})}$$

Sophie's pre-tax withdrawals from her savings during retirement are shown in in figure 2.17. The pre-tax withdrawals are significantly larger than the post-tax withdrawals, though they grow at the same 3% per year. This gap is why tax planning for retirement is so important. We'll quantify how important below when we compute the value of these cash flows, or size of Sophie's nest egg, at the start of her retirement.

Age	30	31	69	70	71	86	86
Period (years)	0	1	39	40	41	56	57
Post-tax Cash flows (\$000)					200	312	321
					$\div(1 - 0.30)$	$\div(1 - 0.30)$	$\div(1 - 0.30)$
Pre-tax Cash flows (\$000)					286	445	459

Figure 2.17: Retirement Consumption with Inflation and Taxes

2.4.2 The Nest Egg

To compute Sophie's nest egg, i.e., the value of her retirement needs at the start of retirement, we can apply the Fundamental Value Relation to the pre-tax withdrawals from her savings. If we do so using her expected return in retirement, 5%, we obtain the following value.

$$Value_{40} = \frac{285,714}{(1 + 0.05)} + \frac{285,714 \times (1 + 0.03)}{(1 + 0.05)^2} + \dots + \frac{285,714 \times (1 + 0.03)^{16}}{(1 + 0.05)^{17}} = \$3,983,829$$

This estimate is 43% larger than the \$2,788,681 nest egg value we computed above when there was no income tax. That is, Sophie needs 43% more savings entering retirement to support the same consumption when she has to pay income taxes on her retirement withdrawals. Ouch!

Investment Earnings Taxes

We shouldn't discount Sophie's retirement cash flows by the expected return on her savings, 5%, because this is a **pre-tax expected return**. The 5% is what she expects to earn on her savings *before* paying taxes on those returns. Her taxes on any investment earnings will depend on precisely how the earnings are generated - dividends from stocks, interest payments on bonds, or capital gains (or price appreciation) on any asset - as well as her overall income level. Again, to keep things simple, we've assumed that Sophie faces a constant 20% tax rate on her investment earnings throughout retirement. So, for each dollar she earns on her savings, she has to pay the government \$0.20. Therefore, her **after-tax expected return** is $0.05 \times (1 - 0.20) = 4\%$. More generally, the after-tax expected return, rt , is computed like so.

$$rt = r \times (1 - TaxRate)$$

Using her 4% after-tax expected return in retirement, the value at the start of Sophie's retirement, 40 years from today, of all her pre-tax withdrawals is

$$Value_{40} = \frac{285,714}{(1 + 0.04)} + \frac{285,714 \times (1 + 0.03)}{(1 + 0.04)^2} + \dots + \frac{285,714 \times (1 + 0.03)^{16}}{(1 + 0.04)^{17}} = \$4,327,778$$

Sophie needs over \$4.3 million at the start of her retirement to cover all of her future expenses and pay her anticipated taxes on income and investment earnings. This estimate is \$1.5 million, or 55%, larger than the estimate in which we ignored the effects of taxes (\$2.8 million).

The size of this difference is why tax planning is so valuable, a point emphasized by the results presented in Table 2.18. The table shows what happens to Sophie's nest egg as the tax rate on income (rows) and investment earnings (columns) change. Even small changes in either tax rate can have a significant impact - hundreds of thousands of dollars - on how much Sophie needs entering retirement. Reductions in either rate, through judicious choices of savings vehicles or investments, can have a profound impact on how much money she'll need for retirement or how much she's able to consume in retirement. The message: understand the tax consequences of your savings and investments or hire someone who does!

		Investment Earnings Tax Rate								
		0.00%	5.00%	10.00%	15.00%	20.00%	25.00%	30.00%	35.00%	40.00%
Income Tax Rate	0.00%	2,788,681	2,846,221	2,905,487	2,966,540	3,029,445	3,094,269	3,161,082	3,229,957	3,300,971
	5.00%	2,935,453	2,996,022	3,058,407	3,122,674	3,188,889	3,257,125	3,327,455	3,399,955	3,474,706
	10.00%	3,098,534	3,162,468	3,228,319	3,296,156	3,366,050	3,438,077	3,512,313	3,588,842	3,667,745
	15.00%	3,280,801	3,348,496	3,418,220	3,490,047	3,564,053	3,640,316	3,718,920	3,799,950	3,883,495
	20.00%	3,485,851	3,557,777	3,631,859	3,708,175	3,786,806	3,867,836	3,951,353	4,037,447	4,126,214
	25.00%	3,718,241	3,794,962	3,873,983	3,955,387	4,039,260	4,125,692	4,214,776	4,306,610	4,401,294
	30.00%	3,983,829	4,066,030	4,150,696	4,237,914	4,327,778	4,420,384	4,515,832	4,614,225	4,715,673
	35.00%	4,290,278	4,378,802	4,469,980	4,563,908	4,660,684	4,760,414	4,863,203	4,969,165	5,078,417
	40.00%	4,647,801	4,743,702	4,842,478	4,944,233	5,049,075	5,157,115	5,268,470	5,383,262	5,501,618

Figure 2.18: The Value of Sophie's Nest Egg at Age 70 Under Different Income and Investment Earnings Tax Scenarios

Growing Annuities

Because Sophie's cash flows in retirement are not constant (Figure 2.17), we can't use our annuity formula shortcut (equation 2.1) to compute the value of these cash flows. However, because the cash flows *grow at a constant rate*, we can use a modified version of our annuity formula.

$$Value_t = \frac{CF_{t+1}}{r - g} \times \left(1 - \left(\frac{1 + r}{1 + g} \right)^{-(T-t)} \right) \quad (2.7)$$

Here, CF_{t+1} is the cash flow one period ahead of when we are valuing the growing annuity, and g is the constant cash flow growth rate. Applying this formula to Sophie's pre-tax retirement needs produces the same estimate of her nest egg as above.

$$Value_{40} = \frac{285,714}{0.04 - 0.03} \times \left(1 - \left(\frac{1 + 0.04}{1 + 0.03} \right)^{-(57-40)} \right) = \$4,327,778$$

Equation 2.7 returns the value of a **growing annuity** as of time t . The rules for using equation 2.7 are similar to those for using equation 2.1 to value non-growing annuities. The cash flows must be (i) equally-spaced in time and (ii) over a finite horizon, except now the cash flows are allowed to vary over time as long as they do so by growing at a constant rate. Note also that to use equation 2.7 the discount rate, r , is must be constant. When the growth rate of the cash flows is zero, equation 2.7 reduces to equation 2.1. No growth is zero growth.

One caveat with growing annuities is that when the discount rate and the cash flow growth rate are equal ($r = g$), we can't use equation 2.7. (We can't divide a number by zero.) Instead, we have to use the following equation.

$$Value_t = \frac{CF_{t+1} \times (T - t)}{1 + r} \quad (2.8)$$

***(Challenging)* A Constant Annuity in Real Terms**

This section shows the relation between financial analysis in nominal and real terms and can be skipped if desired. Figure 2.17 above showed Sophie's retirement needs growing in nominal terms at 3% per year. This is the actual money she withdraws from her savings, pay taxes, and then spends on expenses each year. We can also work in real dollars and real expected returns to get the same results.

For example, Sophie needs \$200,000 at the end of her first year of retirement, 41 years from today. At a 3% rate of inflation, this means she needs $200,000/(1+0.03) = \$194,174.76$ at the start of her retirement, 40 years from today. In other words, what Sophie can buy 40 years from today for \$194,174.76 will cost her $194,174.76 \times (1+0.03) = \$200,000$, 41 years from today; $\$194,174.76 \times (1 + 0.03)^2 = \$206,000$, 42 years from today; $\$194,174.76 \times (1 + 0.03)^3 = \$212,180$, 43 years from today; and so on. Her nominal cash flows - what she actually spends - maintain constant real cash flows of \$194,174.76 - how much stuff she can buy in year 40 dollars. These nominal and real after-tax cash flows are illustrated in Figure 2.19

Age	30	31	69	70	71	86	86
Period (years)	0	1	39	40	41	56	57
Nominal Cash flows (\$000)					200	312	321
					$\div(1+0.03)$	$\div(1+0.03)^{16}$	$\div(1+0.03)^{17}$
Real Cash flows (\$000)					194	194	194

Figure 2.19: Nominal and Real After-Tax Retirement Cash Flows

In real terms, Sophie's retirement needs are a (non-growing) annuity. The cash flows are constant, equally-spaced in time, and end after a finite number of periods (17 years). Therefore, we can use our annuity formula, equation 2.1, with the real cash flow and real discount rate. The real before-tax discount rate in Sophie's case is

$$rr = \frac{1+r}{1+\pi} - 1 = \frac{1+0.05}{1+0.03} - 1 = 1.94\%.$$

Plugging this into our annuity shortcut formula produces the same estimate of our nest egg when we ignored taxes.⁶

$$\frac{194,174.76}{0.0194} \times \left(1 - (1 + 0.0194)^{-(57-40)}\right) = \$2,788,680.53$$

To get the value of our nest egg accounting for both income and investment earnings taxes, we follow a similar recipe using pre-tax cash flows and after-tax investment returns. Start with the first pre-tax cash flow in retirement, $200,000/(1-0.30)=\$285,714$. Deflate this number to year 40 purchasing power by dividing by one plus expected inflation, $285,714/(1+0.03)$

⁶The equivalence of working in nominal and real terms in this setting can be seen from a slight manipulation of the Fundamental Value Relation. The value in year 40 of Sophie's retirement needs expressed in nominal cash flows and the nominal discount rate is given by this expression.

$$Value_{40} = \frac{200}{(1+0.05)} + \frac{200 \times (1+0.03)}{(1+0.05)^2} + \dots + \frac{200 \times (1+0.03)^{16}}{(1+0.05)^{17}}$$

Deflating the \$200,000 in year 41 by 3% produces a real cash flow of \$194,174 in year 40 dollars.

$$Value_{40} = \frac{194 \times (1+0.03)}{(1+0.05)} + \frac{194(1+0.03)^2}{(1+0.05)^2} + \dots + \frac{194 \times (1+0.03)^{17}}{(1+0.05)^{17}}$$

Now recognize that $(1+0.05)/(1+0.03)$ equals one plus the real rate of return, $(1+0.0194)$, and write the value as

$$Value_{40} = \frac{194}{(1+0.0194)} + \frac{194}{(1+0.0194)^2} + \dots + \frac{194}{(1+0.0194)^{17}}$$

This expression is the value of Sophie's retirement needs expressed in real terms - cash flows and discount rate.

= \$277,393. This result means that \$285,714 when Sophie is 71 has the same purchasing power - can buy the same stuff - as \$277,393 when she's 70. Her nominal and real pre-tax cash flows are illustrated in Figure 2.20 and show that the growth of her nominal cash flows, which is equal to the rate of inflation, maintains constant real cash flows or constant purchasing power during her retirement years.

Age	30	31	69	70	71	86	86
Period (years)	0	1	39	40	41	56	57
Nominal Cash flows (\$000)					286	445	458
Real Cash flows (\$000)					$\div(1 + 0.03)$	$\div(1 + 0.03)^{16}$	$\div(1 + 0.03)^{17}$
					277	277	277

Figure 2.20: Nominal and Real After-Tax Retirement Cash Flows

Sophie's after-tax investment return was computed earlier as $5\% \times (1 - 0.20) = 4\%$. However, this is a nominal return, which measures the money she actually receives after paying taxes on her investment earnings. It doesn't measure how much more or less stuff she can buy with the money. Because we are working with real cash flows, we have to use a real discount rate.

$$\text{After-tax real expected return} = \frac{1 + 0.04}{1 + 0.03} - 1 = 0.97\%$$

Using our annuity formula (equation 2.1) produces a nest egg value equal to what we computed above.

$$\frac{277,393}{0.0097} \times \left(1 - (1 + 0.0097)^{-(57-40)}\right) = \$4,327,778.46$$

To finish our discussion here, we should note that deflating the cash flows by inflation resulted in a constant set of real cash flows only because the nominal cash flows were growing at the same rate as inflation, 3%. If the nominal cash flows grew quicker or slower than inflation, deflating by inflation would not result in a constant set of real cash flows, and we would not be able to use our simple annuity formula. If we wanted to use the non-growing annuity formula, we would need to deflate the cash flows and the expected return by the cash flow growth rate. However, doing so means we can no longer interpret the deflated cash flows and expected return as being "real" or corresponding to purchasing power.

2.4.3 The Savings Strategy

Present Value of the Nest Egg

Now that we know Sophie's retirement needs during and at the start of retirement, we can estimate how much she has to save while working. To do so, let's first compute the value of her nest egg as of today, i.e., present value. From our previous calculations, Sophie needs \$4,327,778 at the start of her retirement 40 years from today. The expected return on her savings while working is 10%, higher than the 5% in retirement because she plans on taking a more aggressive investment strategy (e.g., investing in stocks as opposed to bonds). If we discount her nest egg at 10%, we get a present value of

$$\frac{4,327,778}{(1 + 0.10)^{40}} = \$95,621.95.$$

However, this estimate ignores the effect of taxes she has to pay on her investment earnings. Her after-tax expected return is

$$10\% \times (1 - 0.20) = 8\%.$$

Discounting Sophie's nest egg by this rate produces a present value of

$$\frac{4,327,778}{(1 + 0.08)^{40}} = \$199,211.68.$$

If Sophie had \$199,211.68 today and invested this money for 40 years at 10%, she would have \$4,327,778 at the end of the 40 years after paying taxes of 20% each year on her investment earnings. This estimation is illustrated in Figure 2.21.




Figure 2.21: Retirement Savings Timeline - Present Value of Nest Egg Accounting for Inflation and Taxes

Annual Savings While Working

To estimate Sophie’s annual savings previously, we solved the annuity formula (equation 2.1) for the cash flow and arrives at our annuity cash flow formula (equation 2.2). Because she wants her savings to grow by 4% each year, we can’t use these results. Instead, we need to solve the growing annuity formula (equation 2.7) for the cash flow. Doing so yields the following expression.

$$CF_{t+1} = \frac{Value_t \times (r - g)}{\left(1 - \left(\frac{1+r}{1+g}\right)^{-(T-t)}\right)} \quad (2.9)$$

This equation produces the *first* cash flow of the growing annuity. Each cash flow thereafter grows at a constant rate g . In Sophie’s case, her first year’s savings occurring one year from today (period 1, age 31) is

$$CF_{t+1} = \frac{199,212 \times (0.08 - 0.04)}{\left(1 - \left(\frac{1+0.08}{1+0.04}\right)^{-40}\right)} = \$10,229.$$

Her savings the following year at age 32 will be $10,229 \times (1+0.04) = \$10,638$. While this might seem like little in comparison to her nest egg goal of \$4.3 million, her last year of savings at age 40 will be $10,229 \times (1 + 0.04)^{39} = \$47,221$, and she is earning 8% after-taxes every year.

Sophie’s final savings plan is illustrated in Figure 2.22.




Figure 2.22: Retirement Savings Strategy with Taxes & Inflation

2.4.4 Reactions and Comments

Reactions to this and other exercises we’ll carry out in this book often start with “this is totally unrealistic because...”

- “We can’t know the future.” Of course, you can’t. If you could, you wouldn’t need to read this book (or any other book for that matter)! It’s precisely because we can’t predict the future that we *have to* undertake this exercise. Doing so forces us to confront what is, and what is not possible. For example, if you’re imagining a long relaxing retirement, you’d better make sure you have the resources to support it. But, the only way to have any idea of how much you need, when you need it, and what risks you have to take is to perform this exercise.
- “All of the numbers are going to be wrong.” This is a twist on the previous reaction. And, related to the previous response, there is no reason we can’t explore multiple sets of numbers. In other words, when constructing a savings retirement plan, we should have a baseline plan like we constructed above, plus several contingency plans in which we change the assumptions (e.g., investment earnings, time to death, retirement age, taxes, etc.) to understand potential risks. Doing so is easy. It’s the same calculations just with different numbers.
- “It’s too difficult (or complicated).” No it’s not. It just takes effort. Every high school student, and a number of junior high school students, I have taught have not only been able to grasp what we’ve done here but in many instances extend it to incorporate additional aspects of reality. So, “too difficult” is just an excuse for “I don’t want to expend the effort.”

Some of us may think this process is unimportant because they don’t earn much money or are still in school and have no income. First, saving a little over long periods of time can have a big impact because of compounding. For example, saving \$1,200 a year earning 10% per year - the average return on the U.S. stock market since 1927 - will generate \$531,111.06 dollars at the end of 40 years - a little less after-taxes. That’s hundreds of thousands of dollars from saving \$100 a month. Second, not being able to save now or in the near future does not change the importance of creating a savings strategy. It only means that our savings for the next few years will be zero.

2.5 The Value of College

Does it make financial sense to go to college? Tuition and related expenses (room, board, textbooks, etc.) can run over \$100,000 as of 2024. Additionally, we’ll miss out on four years of full-time income while we’re in school. Let’s explore this question using the Wharton School at the University of Pennsylvania (Penn) as a test case and income data as of 2023.

- According to ZipRecruiter, high school graduates in Philadelphia, Pennsylvania, where Penn is located, earn an average annual income of \$45,275. At an average tax-rate of 18.9%, this translates into \$36,718 a year after taxes. If we don't go to college, we will work for 40 years receiving our first paycheck one year from today, and our after-tax income will be assumed to grow by 4% per year thereafter.
- According to Wharton's website, the median first-year annual income after graduating from Wharton's undergraduate program is \$110,000. At an average tax-rate of 26.3%, this translates into \$81,070 a year after taxes. If we go to college, we will work for 36 years receiving our first paycheck one year after we graduate (five years from today), and our after-tax income will grow by 6.7% per year thereafter.
- Tuition and fees for Penn (and Wharton) are \$80,000 per year payable at the start of each year for four years. These costs are assumed to grow at 3% per year.
- Our opportunity cost of capital is 5% per year.

For now, we'll assume we have enough money to pay for college. We'll explore the value of school when we have to finance the costs, i.e., take out a student loan, in the next chapter. We'll also ignore the increasing tax rate as our income grows. While we've seen just how large an impact taxes can have on our income and investment earnings, they'll do little to affect the choice between going or not going to college. At the end of our analysis, we'll explore alternative choices in terms of careers and colleges.

2.5.1 Skip College

If we skip college, the only cash flows relevant for valuation are our after-tax earnings, which are presented in figure 2.23.




Figure 2.23: Skip College Timeline: After-Tax Earnings

Using our Fundamental Value Relation, we can compute the present value of our after-tax earnings by discounting each cash flow and summing.

$$Value_0 = \frac{36,718}{1 + 0.05} + \frac{38,187}{(1 + 0.05)^2} + \dots + \frac{169,504}{(1 + 0.05)^{40}} = \$1,167,764$$

Alternatively, we can recognize these cash flows as a growing annuity, whose present value can be found using equation 2.7.

$$Value_0 = \frac{36,718}{0.05 - 0.04} \times \left(1 - \left(\frac{1 + 0.05}{1 + 0.04} \right)^{-40} \right) = \$1,167,764$$

The present value of our lifetime after-tax earnings is \$1,167,764.

2.5.2 Go to College

If we go to college, we have to pay for school and delay our work income for four years since we won't receive our first paycheck until one year after we graduate. (We have to work *before* getting paid.) Figure 2.24 presents the cash flows for this choice.




Figure 2.24: Go to College Timeline: College Costs and After-Tax Earnings After Graduation

The first four cash flows correspond to our tuition payments to Penn, which are growing at 3% per year. We begin work four years from today, after we graduate, and earn income that grows at 6.7% per year.

As always, we can discount and sum these cash flows to get their value.

$$Value_0 = -80,000 + \frac{-82,400}{1 + 0.05} + \frac{-84,872}{(1 + 0.05)^2} + \dots + \frac{784,540}{(1 + 0.05)^{40}} = \$2,760,232$$

The present value of going to school is \$2,760,232.

Alternatively, we can think of these cash flows in Figure 2.24 as consisting of two growing annuities: (i) the college payments, and (ii) our earnings from work. The present value of the college payments at a 5% opportunity cost is

$$80,000 + \frac{82,400}{(0.05 - 0.03)} \times \left(1 - \left(\frac{(1 + 0.05)}{(1 + 0.03)} \right)^{-3} \right) = \$310,973.$$

The present value of our earnings starts with the value of those earnings as of four years from today. Using our growing annuity result (equation 2.7) yields

$$Value_4 = \frac{81,070}{0.05 - 0.067} \times \left(1 - \left(\frac{1 + 0.05}{1 + 0.067} \right)^{-(40-4)} \right) = \$3,733,068$$

Note, we have only 36 years of earning, instead of 40, because we went to school for four years. Also, \$3,733,068 is the value of our after-tax earnings as of four years from today. Therefore, we need to discount this amount back to today before deducting the present value of the cost of college.

$$Value_0 = \frac{Value_4}{(1 + 0.05)^4} = \frac{\$3,733,068}{(1 + 0.05)^4} = \$3,071,205$$

The difference between the present values of the school costs, \$310,973, and our earnings, \$3,071,205, is \$2,760,232, the same number as above.

2.5.3 The Decision

To summarize,

- If we don't go to college, the present value of our after-tax earnings is \$1,167,764
- If we go to college, the present value of our after-tax earnings less the present value of the cost of college is \$2,760,232.

Going to Wharton is clearly the best decision from a financial perspective, generating an additional $2,760,232 - 1,167,764 = \$1,592,468$ of value for us. In fact, it will always be the best decision as long the income tax rate is less than 64% or Wharton graduates earnings grow more quickly than 2.2% per year or Penn's tuition grows by less than 155% per year.

What this analysis does *not* show is that it is always better to go to college. Wharton's first-year salary for undergraduates is at the very high end of starting salaries for college graduates. Likewise many trades, such as plumbers, electricians, and HVAC (heating, ventilation, and air conditioning) technicians, can earn higher starting salaries and experience higher earnings growth over their career. For example, the average annual starting salary for an undergraduate student graduating with a degree in Sociology from the School of Arts and Sciences at Penn is \$65,833 in 2023, according to Penn's career services. The average salary for an entry-level plumber in Pennsylvania is \$59,320 in 2023 according to ZipRecruiter. At a 20% tax rate, the corresponding after-tax earnings are \$52,666 and \$47,456, respectively.

Assuming both salaries grow at 6% per year and a 5% opportunity cost of capital means the present values of skipping college and going to college are \$2,187,920 and \$1,451,137. In other words, going to college results in $2,187,920 - 1,451,137 = \$736,783.01$ *less* financial value in today's terms. (We should try to replicate these estimates by repeating the process

above with the new salary and growth estimates.) The higher earnings of the college degree in this example cannot compensate for the cost of school and the foregone earnings while in school. Again, this example does *not* suggest that going to college is a bad idea. It only highlights that financial value should be estimated and considered in such an important decision.

2.5.4 The Costs and Benefits of Going to School

Another useful approach to analyzing the choice between going and not going to college examines the **incremental cash flows**. Rather than examining the two sets of cash flows separately and comparing present values, we can examine the difference in the cash flows generated by the two scenarios and estimate the present value of the differences. Figure 2.25 shows the incremental cash flows of going to college and, in doing so, highlights the costs and benefits involved in our choice.

Years	0	1	2	3	4	5	...	40
Go to College	-80.0	-82.4	-84.8	-87.4	0.0	81.1	...	784.5
Skip College	0	36.7	38.2	39.7	41.3	43.0	...	169.5
Difference	-80.0	-119.1	-123.1	-127.1	-41.3	38.1	...	615.0

Figure 2.25: Cash flows for Going to School, Skipping School, and the Difference (\$000s)

The first five cash flows are negative and corresponds to the cost of college, which are substantial. Not only do we have to pay for college, we miss out on income from working. The present value of the costs of going to school is

$$-80,000 - \frac{119,118}{(1+0.05)} - \frac{123,059}{(1+0.05)^2} - \frac{127,132}{(1+0.05)^3} - \frac{41,303}{(1+0.05)^4} = -\$448,865$$

The benefits of college in the form of higher earnings only kick in five years from now. The present value of the benefits of going to school is

$$\underbrace{\left[\frac{38,115}{(1+0.05)} + \frac{41,829}{(1+0.05)^2} + \dots + \frac{615,036}{(1+0.05)^{36}} \right]}_{Value_4} \times \frac{1}{(1+0.05)^4} = \$2,041,333$$

The difference of the present values of costs and benefits is the net present value of going to school relative to skipping school: $2,041,333 - 448,865 = \$1,592,468$, the same number we computed above.

2.6 Problem Solving Tips

Let's summarize and abstract how we tackled the applications in this chapter because we will tackle many of our applications in a similar manner.

1. Draw a timeline as in figure 2.26.




Figure 2.26: Step 1: Draw a Timeline

Time periods are placed above the line, cash flows below. Period “0” corresponds to today or now. Cash outflows are indicated with negative numbers, inflows with positive numbers. When the distinction between inflow and outflow is obvious, we can safely ignore the sign. Note the units of time (e.g., years, months, days) and, when necessary, of cash flows (e.g., \$000s, €million) to the left of the timeline.

In the applications above, we sometimes had several rows of cash flows to denote different costs and benefits (negative and positive cash flows). This level of detail is not only informative but also helps avoid confusion in more complex settings. We can always add or subtract vertically, i.e., cash flows with the same time units.

2. Identify the opportunity cost, r , used to discount or compound cash flows (i.e., the conversion factor used to change the time unit of the cash flows so we can add or subtract cash flows at different points in time). Remember that the opportunity cost must reflect the risk of the cash flows to which they are being applied.
3. Don't forget the implications of taxes. Taxes can affect the cash flows, like an income tax, or they can affect the discount rate, like an investment earnings tax.

For those familiar with spreadsheet programs (e.g., Excel, Google Sheets), timelines should look familiar. They're nothing more than handwritten spreadsheets. Ultimately, timelines play an important practical and conceptual role in solving financial problems. They help us organize and display relevant information that is easily translated into a computational program.

With the relevant cash flow information displayed on our timeline and a discount rate in hand, estimating value and coming to a decision is merely a matter of applying the

Fundamental Value Relation, equation 1.2. Of course, we need to know at what point in time we want a value - today, next year, 10 years from today, etc. But, all of the information we need is neatly organized for us on our timeline. So, the punchline of this short but important section is to reiterate and emphasize what we stated earlier. Get in the habit of drawing a timeline to solve financial problems.

2.7 Key Ideas

We saw in Chapter 1 that you need to know one equation to do most everything in finance, our Fundamental Value Relation.

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

Every equation we used in this chapter is just a special case of this one. Annuities (equation 2.1) and perpetuities (equation 2.3) are the cases where the cash flows are all the same; $CashFlow_1 = CashFlow_2 = CashFlow_3 = \dots$. Growing annuities (equation 2.7) are the cases where the cash flows are all growing at a constant rate; $CashFlow_2 = CashFlow_1 \times (1+g)$, $CashFlow_3 = CashFlow_1 \times (1+g)^2$, \dots . Annuity cash flows (equations 2.2 and 2.9) are just algebraic rearrangements of these equations.

So, in all our applications, we've just been using our Fundamental Value Relation, sometimes explicitly and sometimes implicitly when we use our annuity or perpetuity shortcut formulas. Importantly, our Fundamental Value Relation *always* works, even when the cash flows don't correspond to annuities or perpetuities. Because it's so easy to perform computations in a spreadsheet program, we can always fall back on just discounting (or compounding) the cash flows one at a time and then adding or subtracting.

- We saw in the previous chapter that net present value (NPV) is equal to value less price. An equivalent interpretation of NPV is as the present value of all benefits (cash inflows) minus the present value of all costs (cash outflows). When NPV is positive, the decision is a good one because it creates value. When NPV is negative, the decision is a bad one because it destroys value.
- A useful strategy for tackling problems in finance consists of:
 1. Drawing a timeline to lay out the cash flows. When we create financial models in a spreadsheet, this is exactly what we are doing - identifying all of the relevant cash flows and when they occur.

2. KISS (Keep It Simple Stupid). Consider our retirement savings application. We started with a simple case in which there were no taxes or inflation and Sophie was saving and consuming a constant amount while working and in retirement, respectively. This setting was not particularly realistic, but it let us understand the problem without getting lost in the details. Once we understood the problem, we made things more realistic by adding inflation, taxes, and different investment strategies. Even still, our strategy abstracted from many realities. That's OK. Where to stop adding complexity is up to us and based on our comfort level. The key is to never lose sight of what we're trying to accomplish, which is to make better - not perfect - decisions.

- The present value of a growing annuity at time t is

$$Value_t = \begin{cases} \frac{CashFlow_{t+1}}{r-g} \times \left(1 - \left(\frac{1+r}{1+g}\right)^{-(T-t)}\right), & \text{if } r \neq g \\ \frac{CashFlow_{t+1}}{1+r} \times (T-t), & \text{if } r = g \end{cases}$$

where $CashFlow_{t+1}$ is the cash flow one period ahead, r is the discount rate, g is the cash flow growth rate, and T is time at which the annuity ends (so $T-t$ is the number of periods in the annuity). When $r = g$, we can't use the top expression because we can't divide by zero. Intuitively, what's happening is that the cash flows are growing at the same rate at which we are discounting them so the two effects cancel. Another special case is when the cash flows don't grow, in which case $g = 0$ and the cash flows are all the same, $CashFlow_{t+1} = CashFlow$ for all t . In this case, we get our original annuity result.

$$Value_t = \frac{CashFlow}{r} \times \left(1 - (1+r)^{-(T-t)}\right)$$

- The present value of a growing perpetuity at time t is:

$$Value_t = \frac{CashFlow_{t+1}}{r-g}$$

where $CashFlow_{t+1}$ is the cash flow one period ahead, r is the discount rate, g is the cash flow growth rate. This equation only makes sense if $r - g > 0$. In words, the cash flow growth rate must be strictly less than the discount rate. If not, the present value is infinite.

2.8 Technical Appendix

2.8.1 Derivation of Shortcut Formulas - Growing Annuity and Perpetuity

This appendix derives the shortcut formulas for common cash flow streams, specifically, annuity (equation 2.1) and perpetuity (2.3) and their constant growth versions (equations 2.7 and 2.7). Let's focus on a growing annuity because each result is just a special case.

$$\begin{aligned} & \frac{CF}{1+r} + \frac{CF(1+g)}{(1+r)^2} + \dots + \frac{CF(1+g)^{T-1}}{(1+r)^T} \\ &= \frac{CF}{1+g} \underbrace{\left(\frac{1+g}{1+r} + \frac{(1+g)^2}{(1+r)^2} + \dots + \frac{(1+g)^T}{(1+r)^T} \right)}_{\text{Partial sum of geometric series}} \end{aligned} \quad (2.10)$$

The terms in parentheses form a partial sum of a geometric series which can be expressed as

$$\sum_{t=1}^T a_t = a_1 \left(\frac{1+c^T}{1-c} \right) \quad (2.11)$$

where a_t is the t^{th} term in the sum and c is the common ratio such that

$$a_t = a_1 \times c^{T-1}$$

Note, equation 2.11 only holds true when c is not equal to one.

In our example, $c = (1+g)/(1+r)$ and $a_1 = (1+g)/(1+r)$. Plugging these into the right hand side of equation 2.11 yields

$$\frac{1+g}{1+r} \left(\frac{1 - \left(\frac{1+g}{1+r}\right)^T}{1 - \frac{1+g}{1+r}} \right).$$

Simplifying produces

$$\begin{aligned} \frac{1+g}{1+r} \left(\frac{1 - \left(\frac{1+g}{1+r}\right)^T}{1 - \frac{1+g}{1+r}} \right) &= \frac{1+g}{1+r} \left(\frac{1 - \left(\frac{1+g}{1+r}\right)^T}{\frac{r-g}{1+r}} \right) \\ &= \frac{1+g}{r-g} \left(1 - \left(\frac{1+g}{1+r}\right)^T \right). \end{aligned}$$

Plugging this result into the parenthetical term in equation 2.10 produces the present value of a growing annuity result as of today, time $t = 0$.

$$\frac{CF}{r-g} \left(1 - \left(\frac{1+g}{1+r}\right)^T \right) = \frac{CF}{r-g} \left(1 - \left(\frac{1+r}{1+g}\right)^{-T} \right) \quad (2.12)$$

For the value at an arbitrary time, t , we deduct t from T to get equation 2.7.

The special cases of this result are as follows.

- When the discount rate and growth rate are equal, $r = g$, the common ratio is one and the result in equation 2.11 does not hold. However, looking at equation 2.10, we can see that the parenthetical term is just the sum of T ones which equals T . So, when $r = g$, the present value of a growing annuity is

$$\frac{CF \times T}{1 + g} = \frac{CF \times T}{1 + r}.$$

- When there is no growth, $g = 0$, equation 2.12 reduces to our annuity result, equation 2.1

$$\frac{CF}{r} (1 - (1 + r)^{-T})$$

- When the number of time periods T gets arbitrarily large we have our growing perpetuity result, equation 2.7

$$\lim_{T \rightarrow \infty} \frac{CF}{r - g} \left(1 - \left(\frac{1 + r}{1 + g} \right)^{-T} \right) = \frac{CF}{r - g}$$

as long as the growth rate g is less than the discount rate r ($r > g$).

- Building on the previous result, if the growth rate is zero then we have our perpetuity result, equation 2.3

$$\frac{CF}{r}$$

2.9 Problems

2.1 (*Saving for a purchase*) Larry wants to buy a new Sony Walkman. He also has access to a savings account that pays 10% per year.

Using this information, answer the following questions.

- How much does he have to deposit in his savings account today to be able to buy the Walkman one year from today assuming the price will be \$85?
- How much would Larry have to deposit into the savings account today if he wants to buy the Walkman *two* years from today, again assuming the price will be \$85?

- c. Larry has figured out that he can save \$40 today, and \$40 one year from today. Will he have enough money at the end of two years to purchase the \$85 Walkman if he invests his money in the savings account?
- 2.2 (*Future value from saving*) My daughter has given me \$1 to hold for her. If I invest her money in a bank account paying 2% per annum for one year, how much money will she have at the end of the year?
- 2.3 (*Future value from saving*) A certificate of deposit (CD) is a financial product offered by banks and other financial institutions. It acts like a savings account with restrictions on when you can withdraw your money without incurring a penalty. As a result, it offers higher interest rates than traditional savings accounts in which you can withdraw your money at any time without penalty. If you deposit \$100 into a CD that pays 3% per year, how much will your investment be worth after five years?
- 2.4 (*Valuing a gift*) Your rich friend Dave, who is very trustworthy, has promised to give you \$100,000 one year from today.

Using this information, answer the following questions.

- a. Should the opportunity cost of this money be high or low? Explain.
- b. If the opportunity cost of this money is 5%, what is its present value?
- c. Because you need money today, you have offered to “sell” Dave’s promise of \$100,000 one year from today to another friend, Chuck. How much should you charge Chuck to receive a fair deal?
- d. Dave has come back to you and asked if you would be willing to take payment two years from today. How much would Dave have to pay you two years from today so that you are indifferent between receiving \$100,000 one year from today and waiting an extra year?
- 2.5 (*Affording a loan*) You are looking for financing (i.e., you need money) for a project that you anticipate selling for \$100,000 in five years. Your bank is offering you a 3% annual interest rate on loans. How large of a loan can you take today assuming that you will repay the loan in its entirety from the proceeds of the sale, five years from today?
- 2.6 (*Stock wealth growth*) You have \$1,000 today that you can invest in the stock market. You expect to earn 15% per year. How much money do you expect to have in 20 years?

- 2.7 (*Savings earnings*) If you save \$2,000 per year for thirty years, how much money will you have in 30 years if your expected return is 8% per year and you begin your savings one year from today?
- 2.8 (*Savings earnings*) If you deposit \$20,000 into a savings account today that pays interest at a rate of 6% per annum, how much money will you have in three years? How much total interest will you have earned? How much additional interest is due to compounding (i.e., "interest on interest")?
- 2.9 (*Effect of inflation*) The price of a Genesis GV80 fully loaded (it's a car) is currently \$78,000. Assuming Genesis does nothing to change the car for next year, what do you predict the car will cost if expected inflation is 3.5%?
- 2.10 (*Nominal versus real income*) When setting up your budget for the next several years, you have gathered the following information.
- Your after-tax income receivable one year from today is \$125,000. This income will grow at 4% per year over the next 10 years.
 - The Federal Reserve, the U.S. central bank, forecasts average inflation from 2022 to 2030 to be 3.2% per year.
 - Your opportunity cost of capital is 6% per year.

Using this information, answer the following questions.

- a. What is your nominal income five years from today?
- b. What is your real income in today's dollars five years from today?
- c. What is the real annual growth rate of your income?
- d. What is your real opportunity cost?
- e. What is the present value of your income using nominal cash flows and the nominal discount rate (i.e., opportunity cost of capital)?
- f. What is the present value of your income using real cash flows and the real discount rate?
- g. How do your answers to the previous two questions compare?
- h. If expected inflation were forecast to be 4%, what would your real income be five and 10 years from today? Explain the relation between these numbers.

2.11 (*Valuing an annuity*) A large insurer is offering its clients an annuity that pays \$250,000 per year for 20 years. Your colleague, who has yet to take this finance course, suggests that a fair price for the annuity is \$5 million.

Using this information, answer the following questions.

- a. Under what conditions would your colleague would be correct?
- b. What is a fair price to charge for the annuity if the discount rate for the cash flows is 4% per annum?
- c. The same insurer is offering an alternative annuity with the same features, except the cash flows will grow by 3% per year to maintain pace with expected inflation. What is the most its clients should be willing to pay for this product if the discount rate is 4% per annum?

2.12 (*Really long-term financing*) Costa Rica is considering financing a national military. To do so, they have decided to issue a perpetuity paying 1 billion colones each year starting next year. The current interest rate in Costa Rica is 12%.

Using this information, answer the following questions?

- a. How much money can the government raise today?
- b. In an effort to increase funding, Costa Rica has decided to alter the repayment scheme of its perpetuity so that it grows by 5% per annum after the first year. How much can Costa Rica raise under this new financing scheme?

2.13 (*Value a future annuity*) Steven is considering the purchase of an annuity today for his retirement, which will start 30 years from today. So, the first payment from his annuity would be received 31 years from today. He wants a product that pays \$250,000 in the first year and then grows at 3% per annum to keep up with expected inflation. The annuity will make payments for 20 years. If the opportunity cost facing Steven is 5% per annum, what is the value of the annuity today?

2.14 (*Bank loan repayment schemes*) You are considering taking out a small business loan today for \$250,000. The loan will last for five years and comes with a 7% annual interest rate. The bank has offered you three different types of loans corresponding to different repayment schemes:

- (a) An **amortizing** loan in which you repay the loan in equal annual installments over the five years, beginning at the end of the first year.

- (b) A **coupon** loan in which you pay interest (i.e., 7% of \$250,000) every year for five years starting at the end of the first year, and in the last year you repay the principal (\$250,000).
- (c) A **zero coupon** loan in which you pay the accumulated interest and principal at the end of the five years.

(The last two loans are sometimes referred to as “bullet” loans because of the large principal payment at the end of the loan.)

Compute loan payments for each loan. What might you consider in deciding among the three loans?

- 2.15 (*When to retire*) You currently have \$100,000 and are planning on retiring when your savings reach \$500,000. If the expected return on your savings is 8% per annum, and you don't plan on contributing any additional money towards savings, how long will you have to wait to retire? *Extra credit:* Create a line plot showing the time to retirement as a function of the expected return on your savings.
- 2.16 (*Implied investment returns*) Max is currently ten years old. His father has offered to safeguard his current savings of \$500 until he turns 18. In return, Max has demanded he receive \$1,000 upon turning 18. What is the annual interest implied by Max's and his father's investment scheme?
- 2.17 (*Investment earnings and taxes*) Sophie is ten years old today and has \$1,000 saved. She has asked her father to hold this money until she is 18 years old. In return, she has requested a 10% annual interest rate on her investment.

Using this information, answer the following questions.

- a. How much money will Sophie have in eight years?
- b. Sophie is anticipating an annual inflation rate of 2% per annum. What is her real annual rate of return, and how much purchasing power will she have in eight years?
- c. Sophie's father has decided to tax her earnings on an annual basis at a rate of 20%. He will deduct from her account what is owed in taxes each year. What is Sophie's after-tax interest rate? How much money will she have in eight years?
- d. After thinking about it, Sophie argues that taxing her on unrealized gains is unfair and asks (demands) to be taxed on the aggregate earnings when she receives the cash in eight years. Assuming her father taxes these earnings upon realization

at 20%, how much money will she have in eight years after taxes are deducted, what is her effective annual interest rate under this tax scheme, and how does it compare to the previous tax strategy?

2.18 (*Choosing among pension options*) Three payment options are offered by your pension provider:

- (a) \$70,000 every year starting next year for 20 years.
- (b) A lump sum payment today of \$100,000, and \$60,000 per year starting next year for 20 years.
- (c) \$66,000 per year starting next year for the first 10 years, and \$76,000 per year for the last 10 years.

If your opportunity cost of capital is 8%, which option should you choose?

2.19 (*Repaying an allowance*) Your parents make you the following offer. They will pay you \$12,000 per year starting one year from today for the next 10 years if you agree to pay them back \$15,000 per year for 20 years, starting one year after you receive their last payment.

Using this information, answer the following questions.

- a. Should you take their offer if the opportunity cost of capital is 6% per annum?
- b. At what opportunity cost would you be indifferent between accepting and rejecting your parent's offer?

2.20 (*Valuing an estate*) Your grandfather has set up his estate to distribute annual cash flows in perpetuity according to the following schedule: Compute the value of your

Year	Distribution
1	\$100,000
2	\$200,000
3	\$300,000
4	\$400,000
⋮	⋮

grandfather's estate assuming that the opportunity cost of capital is 6%.

2.21 (*Corporate valuation*) You recently acquired a profitable startup that is expected to generate \$500,000 in profits at the end of the year. Analysts have valued the company at

\$10,000,000. If the company's cost of capital is 9%, what is the constant profit growth rate implied by the analysts' valuation assuming the firm will operate indefinitely?

2.22 (*Getting value from an annuity*) Lois bought an annuity from Rock Solid Insurance for \$4,000,000 when she retired. The annuity is structured to pay her \$400,000 per year until she passes away. How long must Lois live to realize the full value of her investment if her opportunity cost of capital is 9% per annum?

2.23 (*Lottery winnings*) Nelly Farrell of Long Island, New York recently won the \$758.7 million Powerball jackpot. She was given two options to receive her winnings.

- (a) Thirty annual payments, beginning today, of \$25.29 million.
- (b) Lump-sum payout today in the amount \$432,500,936.00.

Nelly's opportunity cost of capital is 5% per annum.

Using this information, answer the following questions.

- a. Ignoring taxes, which option should Nelly choose?
- b. What opportunity cost would make Nelly indifferent between lump sum and the annuity payments?
- c. If Nelly's income is taxed at 36% in the year it is received, which option should she choose?

2.24 (*Professorial endowments*) A grateful former student gave his finance professor \$5 million today. Assume that the annual risk-free rate of return is 5% for the indefinite future.

Using this information, answer the following questions.

- a. What constant amount can the professor draw down each year, beginning immediately, so that he is left with nothing after 30 years?
- b. How would your answer to question 1 change if the professor wanted to ensure that he had \$1 million at the end of 30 years?
- c. What is the maximum, constant amount the professor could draw each year in perpetuity and never run out of money?

2.25 (*Rule of 72 logic*) The "Rule of 72" is a rule of thumb stating that the amount of time it takes to double your investment is approximately equal to 72 divided by the rate of return on your investment. Consider investing \$100 at 4% per annum.

- a. Using the Rule of 72, how long should it take to double your money?
- b. Exactly how long will it take to double your money?
- c. At what interest rate is the Rule of 72 exactly correct?
- d. Graph the relation between interest rate on the x-axis and the difference between the exact amount of time it takes to double your investment and that implied by the Rule of 72. Describe the relation?

2.26 (*Investment banking salaries*) In 2021, the average investment banking salary for a Wharton MBA was \$150,000 per year. Students also received an average signing bonus of \$50,000. Assume a Wharton MBA works for 10 years and invests their money at 6% per annum.

- a. What is the annual salary in their last year of work?
- b. If their annual salary grows at 7%, what is the present value of their total compensation?
- c. What is the future value of their total compensation at the end of 10 years?
- d. What is the present value of the banker's compensation, if in addition to their annual salary they also receive an average annual bonus equal to 125% of their annual compensation?

2.27 (*Deciding to get an MBA*) Venkat is 32 years old. His current annual salary, to be received one year from today, is \$178,000 and is expected to grow at 3% per annum until he retires at age 64. Concerned about his earnings potential, Venkat is considering enrolling next year in the Wharton Executive MBA (WEMBA) program, which costs \$120,000 per year for two years payable at the start of each year. If he enrolls in WEMBA he will earn his current and projected salary for the next three years as he continues working before starting and while enrolled in the program. After graduating, his salary is expected to experience a one-time increase of 25% and then grow by 5% per annum until he retires at age 64. Venkat's opportunity cost of capital is 7% per annum.

Using this information, answer the following questions.

- a. What is the present value of Venkat's lifetime income if he chooses *not* to enroll in the WEMBA program?
- b. What is the present value of Venkat's lifetime income if he chooses to enroll in the WEMBA program?

- c. What is the net present value (NPV) of the WEMBA degree to Venkat?
- d. Unsure of whether he will experience a one-time 25% salary increase following his graduation, what is the smallest one-time increase that would make Venkat still want to pursue the WEMBA degree, assuming the subsequent income growth remained at 5% per annum?
- e. Recompute the value in the previous problem if instead of growing at 5% per year after the one-time increase, Venkat's post-WEMBA salary growth is 3% per year after the one-time increase.
- f. What is the most the WEMBA program could charge Venkat to leave him indifferent between enrolling and not enrolling?

2.28 (*4% rule for retirement spending*) On December 12, 2022, the Wall Street Journal published an article entitled, "The 4% rule for retirement spending makes a comeback," by Anne Tergesen. At the center of the article is the argument that retirees needing to make their money last for 30 years should "spend no more than 4% of their savings in the first year of retirement, and in subsequent years raise those withdrawals to keep pace with inflation."

Claudia estimates she will need to withdraw \$500,000 at the end of her first year of retirement, which she too expects to last for 30 years. Annual inflation during Claudia's retirement is expected to be equal to the long-run historical average of 3%. The expected return on her savings in retirement is 5% per year.

Using this information, answer the following questions.

- a. (4 points) If Claudia grows her consumption in retirement by expected inflation, how much money will she withdraw in her last year of retirement?
- b. (6 points) If Claudia grows her consumption in retirement by expected inflation, how much money does she need at the start of her retirement, one year before her first withdrawal? I.e., how large must her nest egg be going into retirement?
- c. (4 points) The first year withdrawal of \$500,000 represents what percentage of Claudia's nest egg computed in the previous problem? How does it compare to the "4% rule?"
- d. (4 points) To achieve the "4% rule," how much would Claudia have to withdraw in her first year of retirement?

- e. (4 points) Assuming Claudia didn't want to change her first year withdrawal to achieve the "4% rule," how large would her nest egg have to be?
- f. (4 points) Using the nest egg computed in the previous problem and her retirement cash flows, \$500,000 one year from today growing at 3%, what is Claudia's implied expected investment return in retirement? How does this compare to 4%?
- g. (6 points) Construct a two-way data table in which each row represents a different expected return and each column represents a different expected inflation rate. In each cell of the table should be the ratio of the first year retirement withdrawal (\$500,000) to the nest egg at the start of retirement. If you are using Excel and are comfortable with the data table function, your table should look something like the following figure.

		Ratio of first year withdrawal to nest egg									
		Inflation									
		1.0%	2.0%	3.0%	4.0%	5.0%	6.0%	7.0%	8.0%	9.0%	10.0%
Expected Return	1.0%										
	2.0%										
	3.0%										
	4.0%										
	5.0%										
	6.0%										
	7.0%										
	8.0%										
	9.0%										
	10.0%										

Figure 2.27: Two-Way Data Table

If you are not comfortable with the data table function in Excel, construct a 2 x 2 table in which the expected return and inflation vary from 3% to 4%.

What approximate real rate of return is implied by the 4% rule?

2.29 (*Multi-year guaranteed annuities*) A multi-year guaranteed annuity (MYGA) is a financial product offered by insurance companies that functions like a certificate of deposit (CD) offered by banks. A MYGA is purchased with a single up-front deposit that earns a fixed interest rate for a set period of time (e.g., 5 years), at which point the money may be withdrawn or reinvested in a similar product. An important difference between MYGAs and CDs is that interest earned on MYGAs is tax-free until the money is withdrawn when it is taxed as ordinary income. Interest earned on CDs is taxable as ordinary income each year.

Will has \$100,000 to save and is considering the following products.

- (a) 5-year MYGA with an annual rate of 4.15% offered by New York Life. (These products come with fancy names. This one is called the Secure Term MVA II 5 High-Band)
- (b) 5-year CD with an annual interest rate of 4.15% offered by SchoolsFirst Federal Credit Union.

Will's income is taxed at 36%, his opportunity cost is 2.5%, and he plans on leaving his money in whichever product he chooses for the full term to avoid penalties for early withdrawal.

Assuming Will invests in the MYGA, answer the following questions.

- a. How much money will his account show at the end of each year?
- b. How much interest does his money earn each year?
- c. How much tax does he owe each year?
- d. What are the after-tax cash flows each period associated with this investment?
- e. What is the (after-tax) net present value of this investment for Will?
- f. What are the total and annual after-tax rate of returns on this investment?

2.30 Continuing from the previous problem, assuming Will invests in the CD, answer the following questions.

- a. How much money will his account show at the end of each year?
- b. How much interest does his money earn each year?
- c. How much tax does he owe each year?
- d. What are the after-tax cash flows associated with this investment?
- e. What is the (after-tax) net present value of this investment for Will?
- f. What is the annual after-tax rate of return on this investment? (*You will need a computer or calculator to answer this question.*)
- g. What rate must the CD offer so that Will will be indifferent between the MYGA and CD.

2.31 In December of 2023, Shohei Ohtani signed a contract with the Los Angeles Dodgers that would pay him \$70 million a year for ten years beginning at the start of 2024.

However, in order to relieve some of the financial pressure on the team and allows them to hire more star players, Ohtani agreed to defer \$68 million of his compensation each year until after the contract. In other words, Shohei would only receive \$2 million a year from 2024 to 2033, after which he would receive \$68 million dollars a year from 2034 to 2043.

Shohei's personal tax rate is 42% and his opportunity cost is 7% per year.

Using this information, answer the following questions.

- a. Construct a timeline showing Shohei's gross salary, income taxes, and after-tax income assuming he did *not* defer his salary.
- b. What are the present values of Shohei's gross salary, income taxes, and after-tax income assuming he did *not* defer his salary.
- c. Construct a timeline showing Shohei's gross salary, income taxes, and after-tax income assuming he *does* defer his salary.
- d. What are the present values of Shohei's gross salary, income taxes, and after-tax income assuming he *does* defer his salary.
- e. What is the present value of gross income Shohei is sacrificing by deferring his salary?
- f. What is the present value of Shohei's tax savings as a result of deferring his salary?
- g. What is the present value of after-tax income Shohei is sacrificing by deferring his salary?
- h. Assume that Shohei believes that by deferring his salary and hiring additional star players, he can guarantee a World Series victory during his contract which would generate a one-time bonus payment at the end of his contract. How large would this bonus have to be to equate the after-tax value of deferring his salary to the after-tax value of not deferring his salary?

Chapter 3

Borrowing Money

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

When we don't have enough money to buy something, or we want to use our money for another purpose, we'll **finance** the purchase. In other words, we'll borrow money. But, when should we borrow money? When should we pay it back? What are the implications of these choices?

This chapter

- focuses on loans, their use in financing purchases, and the loan interest rate,
- shows what happens to r , the discount rate in the Fundamental Value Relation, when cash flows come and go at a frequency other than yearly (e.g., monthly, quarterly, semi-annually),
- introduces compound interest and how we can earn or pay interest on interest,
- discusses how interest rates change with the maturity of a loan to form the term structure of interest rates, and
- applies our Fundamental Value Relation to answer several questions including:
 - How can we assess college payment plans?

- Should we go to college if we have to take out a student loan, and what are the implications of the loan for our financial well-being after we graduate?
- How do amortizing loans (e.g., mortgage, auto, student) work and how can we select the “best” one?
- When should we refinance or pay back early our mortgage?
- Should we pay cash or borrow money to purchase a car?
- How should we use credit cards and why they are so costly?
- How does an auto lease work and how does it differ from a loan?
- Should we pay back our loans quicker than we’re supposed to?
- What does the federal reserve bank (i.e., “The Fed”) do and how does it impact the economy?

Fundamentally, what’s new in this chapter is that the “ t ” in our Fundamental Value Relation can represent any period of time - year, quarter, month, day, etc. We just need to ensure that the discount rate, r , is measured on the same time scale. Once we have ensured the t and r are measured consistently with one another, we’re just going to see the Fundamental Value Relation (equation 1.2) and its sibling net value (equation 1.3) in action over and over again.

3.1 Loans and Interest Rates

A **loan** is an agreement between two parties (e.g., people, companies, governments) in which one party - the **lender** or **creditor** - gives money to another party - the **borrower** - in exchange for the promise of repayment in the future. For example, a friend might promise to pay us back tomorrow if we give them money today for lunch. While informal, this agreement is a loan between us - the lender - and our friend - the borrower.

Some more formal and common examples of loans include the following.

1. A **student loan** is an agreement between a prospective college student and the government (or possibly a bank) in which the government pays for the student’s college costs, and in return the student promises to pay back the government typically over the 10 years after they graduate.

2. A **home loan**, or **mortgage**, is an agreement between a home buyer and a bank in which the bank gives the home buyer money to help cover the cost of the home, and in return the home buyer promises to pay back the bank over the next 10, 20, or even 30 years.
3. A **credit card** is a loan between a bank and a consumer. The bank gives the consumer money for purchases and the consumer promises to pay back the bank at least a minimum amount every month.

Some other common loans that we as consumers may not think of as loans include:

4. Our **checking** and **savings accounts** at a bank or credit union are loans from us to the financial institution in which we can demand repayment at any time.
5. When we invest money in a **certificate of deposit (CD)** at a financial institution, we are lending money to the institution and are repaid at a future date.
6. When we invest money in **bonds**, we are lending companies, governments, etc. money which gets repaid to us often over time.

These three examples, all of which are discussed in chapter 7, illustrate that we are typically engaged in more loans that we think, and are often borrowers and lenders at the same time.

Loan principal refers to the amount of money borrowed. Any additional money that must be paid back, in addition to the loan principle, is **interest**, which is determined by the **loan interest rate**. For example, if we borrow \$100,000 to go to college at an interest rate of 6%, then our loan principal is \$100,000 and our loan interest rate is 6%. The loan interest rate determines how much extra must be paid back and therefore represents a cost of the loan to the borrower.

While all loans have the basic form of a promise to return money in the future in exchange for money today, loans can differ in many ways. For example, loans can differ in how they must be repaid. Student, home, and auto loans often require constant monthly payments by the borrower over a fixed horizon (e.g., 10 years). Credit cards have no fixed repayment horizon. Borrowers are free to repay as much as they want, when they want, as long as they meet a minimum monthly amount.

Loans can also differ by how they distribute the money. Home loans and term loans to corporations give the homeowner and corporation a lump sum of money at the start of the loan. A student loan distributes money over the course of a students' education, e.g., twice

a year for four years. A credit card offers borrowers access to money up to a limit specific to each borrower at any time.

Some loans can also make certain demands of the borrower. Home loans require borrowers to give the lender the home should they fail to repay the loan. As such, a home loan is said to be **secured** by **collateral** or an asset of value - in this case the house. A credit card has no such requirement; borrowers don't have to promise to give any asset to the bank should they fail to repay the loan much like student loans. As such, credit cards and student loans are said to be **unsecured**. If borrowers default, i.e., fail to pay back the loan, it is easier for lenders to recover their money when the loan is secured. As such, borrowers face lower interest rates on secured loans relative to unsecured loans. This important fact is illustrated in Figure 3.1



Figure 3.1: Average Interest Rates for 30-Year Mortgages and Credit Cards. Source: Federal Reserve Bank of St. Louis

Related, many corporate loans place explicit restrictions, called **covenants**, on the actions and financial conditions of borrowers that if not satisfied can lead to the lender demanding immediate repayment of the loan. For example, most corporate borrowers must agree to not sell certain assets, not make excessive dividend payments to shareholders, and maintain a certain degree of financial health to insure the lender that they will be repaid.

We'll discuss some of these dimensions in more detail in later chapters. In this chapter, the focus will be on evaluating loans for personal use. Specifically, we want to understand how to apply our Fundamental Value Relation for decision making in the context of borrowing

money when repayment can occur at an arbitrary frequency (e.g., monthly, quarterly)

3.2 Compounding

Chapter 2 focused on examples in which the cash flows arrived and departed yearly or at an annual frequency. We also assumed the opportunity cost, r , was what we expected to earn over a year's time on investments with risks similar to the cash flows we were examining. Now, we're going dig into the computation of r so we can more confidently engage with and more accurately assess financial products that don't rely on a yearly time frame for their cash flows and discount rates. Let's start with something simple and familiar.

Imagine that our bank offers a savings account with an annual interest rate of 10%. If the bank pays us interest once a year, then \$100 in savings will lead to $100 \times (1 + 0.10) = \110 at the end of the year. That is, \$110 is the future value of \$100 one year from today at a 10% annual interest rate. And, \$100 is the present value of \$110 one year from today at a 10% annual interest rate. This computation is illustrated in Figure 3.2 and represents nothing new relative to the previous chapter.




Figure 3.2: Annual Compounding

Because the bank is only paying us interest once a year, the 10% interest rate is often referred to as an **Annual Percentage Rate** or **APR**, and the interest earned during the year is referred to as **simple interest**. But, what if the bank decides to pay us interest twice a year, once every six months? This scenario is illustrated in Figure 3.3, which indicates each period corresponds to six months. The interest earned during each period is dictated by the **periodic interest rate**, which we compute by dividing the APR by the number of periods in a year in which interest is paid. In this example, the periodic interest rate is $10\%/2 = 5\%$.

After the first period, six months, we'll have $100 \times (1 + 0.05) = \105 in our account. After the second period, one year from today, we'll have $105 \times (1 + 0.05) = \110.25 in our account, \$0.25 more than what we had when interest was only paid once a year. This increase is a result of interest earning interest in the second half of the year. The \$5 of interest we earn in the first half of the year is earning 5% interest in the second half of the year, which




Figure 3.3: Semi-annual Compounding - Periodic Rate

amounts to $5 \times 0.05 = \$0.25$. This process of interest earning interest is not new. It's simply compounding, and the interest earned as a result of compounding is referred to as **compound interest**. In this example, interest is compounded semi-annually or twice per year.

Looking at Figure 3.3, we can compute the future value of our money in just one step by utilizing the periodic interest rate and the number of compounding periods in the year.

$$100 \times \left(1 + \frac{0.10}{2}\right)^2 = \$110.25$$

Notice in this expression that our discount rate, $0.10/2$, and time, the exponent "2" are measured in the same time units. The discount rate, $0.10/2$, is a six month periodic interest rate and the exponent "2" is the number of six month periods. The broader lesson from this example is that to move money forward or backward in time, i.e., convert time units, we have to ensure that the discount rate and time are consistent with one another. If we have a six month discount rate, we must measure time in six month increments. If we have a quarterly discount rate, we must measure time in quarters. And so on.

Consider our simple savings example again measuring time in years. We need a discount rate that measures our opportunity cost or expected return over a year. We might think to use the APR introduced above, but that would be a mistake. The APR only measures simple interest, which ignores the effects of compounding. But, the money we earn includes the effects of compounding. In our semi-annual compounding example, we earned \$10.25 in interest, not \$10.00. So, if the interest is compounded at a frequency other than annual, using the APR to discount cash flows is a mistake. Instead, we need to estimate an **effective annual rate** or **EAR**.

The EAR measures how much interest is earned in a year accounting for compounding. From our previous calculation, we know that with semi-annual compounding, we'll earn $110.25 - 100 = \$10.25$ over the year. In other words, the EAR is $10.25/100 = 10.25\%$. However, we can compute the EAR directly like so.

$$EAR = \left(1 + \frac{0.10}{2}\right)^2 - 1 = 10.25\%$$

So, while the APR is 10%. With semi-annual compounding, we are actually earning 10.25% a year. Figure 3.4 illustrates the use of the EAR to compute the future value of our money. Again, its important to emphasize that to move money forward or backward in time, i.e.,



Figure 3.4: Semi-annual Compounding - Effective Annual Rate

convert time units, we must ensure that the discount rate, r , measures the expected return over the same time frame represented by the exponents in our Fundamental Value Relation.

To reinforce our understanding, and highlight a pattern, let's see what happens with monthly compounded interest. Figure 3.5 shows the calculation of interest when we measure time in periods (i.e., months). Our periodic interest rate is $0.10/12 = 0.0083$, which measures how much interest is earned each month. This small amount of interest, in addition to our initial \$100, earns interest every month, resulting in a year-end balance in our account of

$$100 \times \left(1 + \frac{0.10}{12}\right)^{12} = \$110.47.$$

The periodic interest rate, $0.10/12$, measures interest earned over a month, the same unit represented by the exponent, "12."




Figure 3.5: Monthly Compounding: Using the Periodic Rate and Measuring Time in Periods

Figure 3.6 shows the future value calculation when we measure time in years. Our effective annual rate is $(1 + 0.10/12)^{12} - 1 = 10.47\%$. So, over one year, our \$100 will become

$$100 \times (1 + 0.1047)^1 = \$110.47,$$

which is exactly what we computed using the periodic interest rate and measuring time in months. Again, notice the consistency between our discount rate, 10.47%, that measures

2. We introduced three types of interest rates.

- (a) The **annual percentage rate (APR)** measures **simple interest**, which is the interest earned in a year ignoring compounding. The APR is *not* a discount rate! We cannot use the APR to discount cash flows because the APR does not tell us how much interest we actually earn or owe over a period. It is best to think of the APR as an interest rate quote, from which we can derive a discount rate.¹
- (b) The **periodic interest rate**, which we'll denote i , measures the interest earned in a compounding period (e.g., quarter, month, day).

$$i = \frac{APR}{k}. \quad (3.1)$$

The periodic interest rate *is* a discount rate and is used to discount periodic cash flows when measuring time in *periods* (e.g., quarterly, monthly, daily).

- (c) The **effective annual rate (EAR)**, denoted r , measures the interest earned in a year accounting for compounding. Mathematically,

$$r = \left(1 + \frac{APR}{k}\right)^k - 1 = (1 + i)^k - 1 \quad (3.2)$$

The EAR is sometimes referred to as the **annual percentage yield** or **APY**. The EAR is a discount rate and can be used to discount cash flows when time is measured in *years*. The EAR can also be used to recover the periodic discount rate, i ,

$$i = (1 + r)^{1/k} - 1, \quad (3.3)$$

and by extension the APR which equals $i \times k$.

Let's look at some applications to both answer important questions and reinforce our understanding of these concepts.

3.3 College Payment Plans

The University of Pennsylvania (Penn) offers students a payment plan that spreads their tuition expenses over 48 months. Assuming the cost of school doesn't change, and taking a little liberty with the exact nature of the plan, the tuition costs are $320,000 \div 48 = \$6,666.67$




Figure 3.7: College Tuition Payment Plan

per month beginning at the start of school. The timeline of tuition costs under the plan is illustrated in figure 3.7.

If our bank is offering a savings account paying a 5% APR with monthly compounding, how much money do we need at the start of school to ensure we can make all 48 payments? To answer this question, we need to compute the present values of the tuition payments and sum. There are several ways we can do this.

3.3.1 Measuring Time in Periods

Let's starting by computing a monthly periodic interest rate to line up with the monthly cash flows. Using equation 3.1, the monthly periodic interest rate is

$$i = \frac{APR}{k} = \frac{0.05}{12} = 0.4167\%.$$

We can use our Fundamental Value Relation, measuring time in periods (months) and discounting the monthly cash flows with the monthly periodic discount rate.

$$Value_0 = 6,666.67 + \frac{6,666.67}{1 + 0.004} + \frac{6,666.67}{(1 + 0.004)^2} + \dots + \frac{6,666.67}{(1 + 0.004)^{47}} = \$290,692.57$$

Because the first payment is made today, it is not discounted. The exponents in the denominators range from 1 to 47 and correspond to the number of months in the future the payment must be made.

We can also recognize that the second through 48th payments constitute an annuity. Using our annuity formula (equation 2.1) with the periodic interest rate and time measured in months produces the same value.

$$Value_0 = 6,666.67 + \frac{6,666.67}{0.004} \times (1 - (1 + 0.004)^{-47}) = \$290,692.57$$

¹The only time one can use the APR to discount cash flows is when the compounding frequency is annual, in which case the APR, the periodic interest rate, and the EAR are all the same. APRs seen in practice can also include fees associated with the loan. Make sure to understand exactly what APR is measuring.

3.3.2 Measuring Time in Years

We can also compute the present value using the EAR and measuring time in years. The EAR (r) is

$$r = \left(1 + \frac{APR}{k}\right)^k - 1 = \left(1 + \frac{0.05}{12}\right)^{12} - 1 = 5.1162\%$$

The present value calculation looks as follows.

$$\begin{aligned} Value_0 &= 6,666.67 + \frac{6,666.67}{(1 + 0.0512)^{1/12}} + \frac{6,666.67}{(1 + 0.0512)^{2/12}} + \dots + \frac{6,666.67}{(1 + 0.0512)^{47/12}} \\ &= \$290,692.57 \end{aligned}$$

Notice that the fractional exponents ($1/12$, $2/12$, ..., $47/12$) are the ratio of the month in which the payment must be made divided by 12. Thus, the exponents measure time in years. A limitation of measuring time in years in this example is that we can't use our annuity formula, in which T measures the number of cash flows and therefore must correspond with the number of periods.

The \$290,692.57 present value of these tuition payments is less than the present value of the four, \$80,000 annual tuition payments computed in the chapter 2) at the same 5% annual cost of capital - \$297,859.84. There are two reasons for this difference. First, the payments here are spread out over a longer horizon - (almost) four years instead of three. Delaying payments further into the future makes them less costly to us because of the time value of money. Second, we're earning more money on our savings because of the monthly compounded interest. Thus, we need less money today to cover our school costs.

Notice that the university, by dividing the total cost of college by the number of payments ($\$320,000 \div 48$), is giving us a loan with a zero interest rate. Another example of a zero interest rate loan as of 2022 is Apple's iPhone Payments plan which requires 24 monthly payments at a 0% APR. A \$1,200 iPhone can be purchased either by paying \$1,200 today, or 24 monthly payments of $1,200 \div 24 = \$50$. All else equal, we should take advantage of interest free loans because the money we would have spent purchasing the asset (iPhone, college education) can instead be invested. However, we have to be careful with **no interest rate loans**. These loans can include hidden fees and costs so read the fine print. Some loans have zero interest only for a short time.²

²For example, the Citi Diamond Preferred Credit Card has no annual fee and 0% APR for 21 months on balance transfers and 12 months on Purchases. After these "teaser" periods end, the regular APR varies from 13.7% to 23.7% - quite hefty.

3.4 Paying for College with Student Loans

In chapter 2, we asked whether it made financial sense to go to college assuming we had the money today to cover our education expenses. Let's revisit that question assuming we have to borrow the money to pay for school. Our key assumptions were as follows.

- If we don't go to college, we'll earn \$36,718 after taxes at the end of this year. These wages will grow by 4% per year and we'll work for 40 years.
- If we go to Wharton, we'll earn \$81,070 after-taxes at the end of our first year after graduating. We'll work for 36 years during which time our wages will increase by 6.7% per year.
- Tuition and fees for Penn are \$80,000 per year payable at the start of each year for four years. These costs are assumed to grow at 3% per year.
- Our opportunity cost of capital is 5% per year.

Additionally, let's assume the loan APR is 7% with annual compounding and payments, and we don't start repaying the loan until one year after we graduate.³ We have 10 years to repay the loan starting one year after we graduate, but interest accrues on the loan while we're in school. This means the loan balance will grow each year while we're in school because we are not making any payments.

In the previous chapter, we showed the following.

1. If we skip school, the present value of our after-tax earnings is \$1,167,764.
2. If we go to school and pay for it from money we have (e.g., savings, parents), the present value of our after-tax earnings accounting net of the school costs is \$2,760,232.

(For details, see 2.) So the decision to go to Wharton was an easy one because doing so created $2,760,232 - 1,167,764 = \$1,592,468$ of financial value for us.

Now let's explore the value of going to school when we borrow money to pay for it. We might do so if we don't have the money to pay ourselves or the loan terms are particularly attractive. (We'll explain precisely what attractive loan terms are below.)

³Interest on most student loans is compounded monthly to coincide with monthly payments. We're abstracting from that detail to focus on the important issues.

3.4.1 Borrowing to Go to School

A loan for school is referred to as a student loan because it is the student taking out the loan and receiving the money. Student loans are examples of **amortizing loans**, which are loans in which interest and principal are paid gradually over the loan term. Auto loans and home loans (a.k.a., mortgages) are other examples of amortizing loans, and like student loans, borrowers pay the same amount every period over the life of the loan.

Figure 3.8 details the cash flows of going to school by borrowing money. Our school costs in the first four years are covered by our student loan distributions. So, the net cash flows are zero - we don't pay or receive any money. One year after we graduate, five years from today, we get our first paycheck for \$87,400, which grows by 6.7% per year thereafter, and we have to start repaying the loan. Ten years later, in year 14, we make our last loan payment and the only remaining cash flows are our after-tax earnings until we retire 40 years from today.

Years	0	1	2	3	4	5	14	15	40
School costs	-80.0	-82.4	-84.9	-87.4					
Loan dists.	80.0	82.4	84.9	87.4					
Loan pmts.						-56.4	-56.4		
Earnings						81.1	145.3	155.1	784.5
Sum	0	0	0	0	0	24.6	88.9	155.1	784.5

Figure 3.8: Financing College Cash Flows (\$000s)

To determine our loan payments, we need to first figure out how much we will owe - i.e., our loan principal - when we graduate. Put differently, we need to find the future value as of year four of each of the loan distributions and then sum. Careful! To find the future value of the loan distributions, we need to use the loan interest rate (7%), which measures the rate at which interest accrues or accumulates on the loan.

$$\begin{aligned}
 LoanPrincipal_4 &= 80,000(1 + 0.07)^4 + 82,400(1 + 0.07)^3 \\
 &\quad + 84,872(1 + 0.07)^2 + 87,418(1 + 0.07) \\
 &= \$396,515
 \end{aligned}$$

To determine our annual payments, we use the annuity cash flow formula (equation 2.2) again with the loan interest rate.

$$LoanPayment = \frac{LoanPrincipal \times r}{1 - (1 + r)^{-T}} = \frac{396,515 \times 0.07}{1 - (1 + 0.07)^{-10}} = \$56,455$$

We will have to pay \$56,455 each year for ten years starting five years from today - years 5 through 14 - in order to pay off our student loan.

To value going to school by borrowing money, we use the fundamental value relation with the cash flows from Figure 3.8 and our 5% opportunity cost of capital.

$$\begin{aligned} Value_0 &= \frac{24,615}{(1 + 0.05)^5} + \dots + \frac{88,870}{(1 + 0.05)^{14}} + \frac{155,062}{(1 + 0.05)^{15}} + \dots + \frac{784,540}{(1 + 0.05)^{40}} \\ &= \$2,712,565 \end{aligned}$$

Compared to not going to school, going to Wharton on student loans creates an additional $2,712,565 - 1,167,764 = \$1,544,801$ of value.

It seems like an obvious decision to go to Wharton as opposed to not going to school even if we have to borrow. We create more value by going to Wharton than not going to college. However, our analysis is much more than one number. The timeline in Figure 3.8 shows how much money we'll have to live on each year after we graduate. That first year after graduation, five years from today, our after-tax earnings are \$81,070 and our loan payment is \$56,455. This leaves us with \$24,615 on which to live. This is an extremely tight budget pretty much anywhere in the U.S. as of 2023.

Ultimately, the decision to go to college is more than just a financial one. However, by using the tools here, we can make a more informed decision and mitigate the chances of falling into financial straights because of a bad decision.

3.4.2 The Cost of Borrowing Money to Go to School

As we mentioned above and showed in chapter 2, the present value of going to college and paying for it with our own money created \$2,760,232 of value in today's dollars. Going to college and having to finance it created \$2,712,565 of value in today's dollars. The difference, $2,760,232 - 2,712,565 = \$47,667$, suggests that if we have the money to pay for school, we should use it to do so and avoid borrowing money. It's costly.

But, from where does this cost come and is it always positive? In other words, should we always pay for college with our savings - assuming we're able - or are there times when we might want to borrow money and continue saving? This is actually an example of a much broader question: Should we ever borrow money while we have savings? We'll answer this question here in the context of our student loan but revisit it below in the context of a home mortgage.

The cost of the mortgage comes from the difference between the loan rate, 7%, and our opportunity cost, 5%. So, any money we save earns less than what we pay on the loan. This difference is what makes borrowing money costly and less attractive when we have other money at our disposal. To be precise, let's discount the benefits and costs of the loan to get its net present value. The present value (PV) of the loan benefits are the sum of the school payments discounted by our opportunity cost.

$$\text{PV Loan Benefits} = 80,000 + \frac{82,400}{(1 + 0.05)} + \frac{84,872}{(1 + 0.05)^2} + \frac{87,418}{(1 + 0.05)^3} = \$310,972$$

The present value of the loan costs are the sum of the future loan payments discounted by our opportunity cost.

$$\text{PV Loan Costs} = \frac{56,455}{(1 + 0.05)^5} + \frac{56,455}{(1 + 0.05)^6} + \dots + \frac{56,455}{(1 + 0.05)^{14}} = \$358,640$$

The difference, i.e., net present value, is -\$47,667, the same number we computed above as the difference in values between going to school and paying for it with savings versus borrowing.

When the cost of borrowing is less than our opportunity cost, it becomes advantageous to borrow money. Consider a situation in which the loan interest rate is 3% and our opportunity cost remained at 5%. The cash flows of going to college and borrowing money are presented in figure 3.9.⁴

The value of going to Wharton and borrowing money at 3% to pay for school is

$$\begin{aligned} \text{Value}_0 &= \frac{38,848}{(1 + 0.05)^5} + \dots + \frac{103,101}{(1 + 0.05)^{14}} + \frac{155,062}{(1 + 0.05)^{15}} + \dots + \frac{784,540}{(1 + 0.05)^{40}} \\ &= \$2,802,981. \end{aligned}$$

This value is greater than the value of going to Wharton and paying for it with our own money, \$2,760,232. Borrowing money creates 2,802,981. - 2,760,232 = \$42,749 of value. How? Because for every dollar we borrow, we pay \$0.03. For every dollar we save, we earn

⁴To get the new monthly payments, we followed the same approach as before. The loan principal in year 4 is

$$80,000(1 + 0.03)^4 + 82,400(1 + 0.03)^3 + 84,872(1 + 0.03)^2 + 87,418(1 + 0.03)^1 = \$360,163$$

. The annual payment is found with the annuity cash flow formula (equation ??).

$$\frac{360,163 \times 0.03}{(1 - (1 + 0.03)^{-10})} = \$42,222$$

Years	0	1	2	3	4	5	14	15	40
School costs	-80.0	-82.4	-84.9	-87.4					
Loan dists.	80.0	82.4	84.9	87.4					
Loan pmts.						-42.2	-42.2		
Earnings						81.0	145.3	155.1	784.5
Sum	0	0	0	0	0	38.8	103.1	155.1	784.5

Figure 3.9: Financing College Cash Flows at 3% Loan Rate (\$000s)

\$0.05. So rather than taking money out of our savings that's earning 5%, let take someone else's money that costs 3%.

The important lesson here is that borrowing money is “costly” when the loan interest rate is greater than our opportunity cost - a lesson that applies to *all* forms of borrowing, not just student loans. When possible, we should avoid debt whose interest rate is greater than our opportunity cost. When we can't avoid borrowing at a relatively high rate because we don't have enough money for something we need - education, car, house, etc. - then we should try to pay down the debt as quickly as possible if the borrowing cost (loan interest rate) is greater than our opportunity cost.

We might think: “let's save in the stock market! That surely has a higher return than the interest rate on a student loan.” Historically, this is true, but only *on average*. The problem with this logic is that we have to pay back our loans. The stock market does *not* have to return 11% every year and, in fact, it rarely does. Stock market returns are highly volatile, varying over a wide range from year to year. Between 1928 and 2023, the annual return to the S&P500 index, a portfolio of the 500 largest publicly traded companies in the U.S., varied between -44% and 50%. So, using the average return on the stock market as the opportunity cost for our loan is violating a principle we laid out in chapter one: the opportunity cost for a set of cash flows reflects the risk of those cash flows. Because we have to pay back the loan, our opportunity cost should be one associated with a very safe, if not risk-free, investment.

For most of us, our opportunity cost will almost always be less than our borrowing cost. The interest rate on our student loan, home loan, car loan, etc. will almost always be higher than the interest rate on a safe investment. This is to say that most people borrow at a higher interest rate than the one at which they lend. However, there are instances where this isn't true. We'll discuss some below.

3.5 Financing a Home

One of the most significant financial decisions people make is purchasing a home. Home's are expensive to buy, expensive to maintain, and highly **illiquid** or difficult and costly to convert to cash. They can be a source of great pain and great pleasure. In the U.S., the large majority of homes are purchased with borrowed money, i.e., a home loan or mortgage. Let's explore mortgages more deeply with an example.

Imagine we found a home we want to buy and that costs \$625,000. To pay for it, we'll put down 20%, which means we'll write a check for $625,000 \times 0.20 = \$125,000$, and borrow the rest, \$500,000.⁵ Now we have to choose a mortgage product, and there are lots. Most banks and credit unions offer a large variety of mortgage products differing in several ways. Borrowers can choose a loan term or how long the loan lasts, which typically varies among 10, 15, 20, and 30 years. Borrowers can also choose a fixed-rate mortgage in which the loan payments are constant over the life of the loan, or an **adjustable rate mortgage** (ARM) in which the loan payments can change from year to year. Finally, borrowers can choose to pay upfront fees or **points** in exchange for a lower loan interest rate. How to approach these decisions is what motivates our this section. To illustrate concepts and get comfortable with mortgages, let's use the most popular mortgage in the U.S., a fixed-rate, 30-year mortgage with no points and a 3% APR.

3.5.1 Mortgage Mechanics

The timeline for our mortgage is in Figure 3.10. We receive \$500,000 today to purchase the home in exchange for 360 monthly payments of some amount, CF . According to our fundamental valuation relation, the sum of the present values of the cash flows should equal \$500,000, the principal of the loan.




Figure 3.10: Mortgage

Because the cash flows are the same, equally spaced out in time, and for a finite amount of time, they correspond to an annuity. Let's use equation 2.2 to compute the annuity cash

⁵Paying cash for 20% or more of the value of the home (i.e., borrowing 80% or less) typically ensures we don't have to buy private mortgage insurance (PMI), which protects the lender in case the borrower defaults repay the loan.

flow, which in this case corresponds to the monthly mortgage payment. To use the annuity formula, we will have to measure time in periods (months) and use the periodic interest rate, which is $0.03/12 = 0.0025$.

$$CF = \frac{Value_0 \times i}{1 - (1 + i)^{-T}} = \frac{500,000 \times 0.0025}{1 - (1 + 0.0025)^{-360}} = \$2,108.02$$

The $Value_0$ of the annuity is how much we borrowed - the loan principal. The discount rate is the periodic interest rate. And, the number of periods is the number of months, $30 \times 12 = 360$. By paying \$2,108.02 per month for 360 months, we will pay all of the interest and principal on the loan.

Amortization Table

Accompanying most mortgages is an amortization table similar to what we discussed in chapter 2. Table 2 shows the top and bottom of the amortization table for our mortgage. Ignoring period 0, there are 360 rows - one for each payment. Each period is one month in duration. An amortization table is nothing more than a timeline oriented vertically instead of horizontally.

To keep with convention, we start the loan today, the end of period 0, with a balance of \$500,000. But for an arbitrarily small amount of time, there is no difference between the end of a month and the start of the next month. Think of the end of the month as 11:59:59 PM on the last day of the month, and the start of the next month at 12:00:00 AM - a one second difference. So, the loan balance at the start of month 1 is the same as the loan balance at the end of month 0, \$500,000, and similarly for all subsequent months.

Period	Start Balance	Interest	Monthly Payment	Principal Reduction	Ending Balance
0					500,000.00
1	500,000.00	1,250.00	2,108.02	858.02	499,141.98
2	499,141.98	1,247.85	2,108.02	860.17	498,281.81
⋮	⋮	⋮	⋮	⋮	⋮
359	4,200.28	10.50	2,108.02	2,097.52	2,102.76
360	2,102.76	5.26	2,108.02	2,102.76	0.00

Table 2: Mortgage Amortization Table

During each month, interest accumulates or **accrues** at a rate equal to the monthly periodic rate (0.25%). To determine how much interest accrues in any month, we multiply what we owe at the start of the month by the periodic interest rate. For example, the

interest owed at the end of the first month is $500,000 \times 0.0025 = \$1,250$. When we make our mortgage payment, part of the payment goes towards paying the interest that accrued over the month, the remainder goes towards paying down the principal. From our first payment of \$2,108.02, \$1,250 pays the interest that accumulated over the month, and $2,108.02 - 1,250.00 = \$858.02$ pays down the principal. The loan balance at the end of the first month is the starting balance minus the principal reduction or $500,000 - 858.02 = \$499,141.98$.

This process continues for another 359 months. The last payment is exactly equal to the starting balance (\$2,102.76) plus the accrued interest (\$5.26), \$2,108.02. After 360 months, nothing is owed on the loan so no more interest can accumulate and we are done paying off our mortgage.

While we can use the amortization table - which is easy to construct in a spreadsheet - to look up any information about our loan, it's useful to understand how to directly compute some important quantities.

Outstanding Loan Balance vs. Loan Value

To determine the outstanding balance of our loan at any point in time, we compute the present value of the remaining payments discounted by the loan's periodic interest rate. For example, to find the outstanding balance after five years, or the 60th payment, we recognize that there are $360 - 60 = 300$ remaining payments of \$2,108.02. The present value of these payments at the periodic loan interest rate is

$$Principal_{60} = \frac{2,108.02}{(1 + 0.0025)} + \frac{2,108.02}{(1 + 0.0025)^2} + \dots + \frac{2,108.02}{(1 + 0.0025)^{300}} = \$444,531.82.$$

Because the loan payments are an annuity, we can also use equation 2.1.

$$Principal_{60} = \frac{2,108.02}{0.0025} \times (1 - (1 + 0.0025)^{-(360-60)}) = \$444,531.82. \quad (3.4)$$

It's important to emphasize that by using the loan interest rate to discount the loan payments, as opposed to our opportunity cost, we are computing the loan principal - what we owe to the bank - not the *value* of the loan. Generally, we borrow at a higher interest rate than the interest rate at which we can lend. So, the interest rate on our mortgage will typically be higher than our opportunity cost, at least at the start of the mortgage. As a result, the value of the mortgage will be greater than the principal or the amount of money the bank gives us. Put simply, mortgages are costly for us.

Take our example in which the loan starts by giving us \$500,000 at a 3% APR to buy the house. Our opportunity cost must reflect the risk of the cash flows, which are relatively safe

because we have to pay back the loan. So, assume we can earn 2% a year on safe investments, or $(1 + 0.02)^{1/12} - 1 = 0.1652\%$ per month.⁶ The *value* of the loan is the sum of the loan payments discounted at our opportunity cost.

$$\frac{2,108.02}{0.001652} \times (1 - (1 + 0.001652)^{-360}) = \$570,321.73$$

The value in today's dollars of all the loan payments we'll have to make is \$570,321.73. But the bank is only giving us \$500,000 - the loan principal. So, this loan is costing us $572,322 - 500,000 = \$72,322$ in today's terms. If we had enough savings to buy the house with all cash - no loan - we should probably do so, in which case we avoid this cost.⁷

The bank, on the other hand, gets to lend at a higher interest rate than the rate at which it borrows. Assume the bank's opportunity cost, also known as a **cost of capital**, is also 2%. We know from the previous calculation that the present value of all the loan payments at an annual discount rate of 2% is \$570,322. This is the value in today's dollars of what the bank is receiving. Because it is only paying \$500,000, the bank is making a gain of \$70,322 in today's dollars. In practice, the bank's opportunity cost would be even lower than the 2% meaning they're making even more money on the loan than what it's costing the borrower.

Payment Composition

The amortization table in Table 2 shows that each monthly mortgage payment consists of two parts: interest and principal reduction. While the payments are constant, the amount, and therefore the proportions, of each payment used to pay interest and reduce principal change every month. This too can be seen in the amortization table.

Consider the 100th mortgage payment. What fraction of this payment covers the interest expense, and what fraction covers principal reduction? The loan principal *before* the payment is made equals

$$Principal_{99} = \frac{2,108.02}{0.0025} \times (1 - (1 + 0.0025)^{-(360-99)}) = \$403,756.28$$

⁶The 2% is an effective annual rate or what we actually earn in a year. To get the periodic rate, we need to find the monthly rate that when compounded 12 times equals 2%, $(1 + 0.001652)^{12} - 1 = 2.0\%$.

⁷We might think we could just take out the loan and put our savings in the stock market for 30 years knowing we'll probably make more than 3% per year on average. This argument makes several mistakes. First, the stock market is much riskier than our mortgage payments and not an appropriate opportunity cost. Second, if we need to use our savings to pay off the mortgage, the variation in stock returns could really haunt us. In down markets we might not have enough to make the mortgage payments, or we may run out of savings before the end of our mortgage. Finally, stock investments don't become less risky as we hold them longer, they become riskier.

The interest accruing in the 100th month is $403,756.28 \times 0.0025 = \$1,009.39$. Therefore, the principal reduction must be the size of the mortgage payment minus this interest expense, or $2,108.02 - 1,009.39 = \$1,098.63$. So, $1,009.39 \div 2,108.02 = 47.9\%$ of our 100th mortgage payment is going towards paying interest, and $1,098.63 \div 2,108.02 = 52.1\%$ is going towards principal reduction.

At the start of the mortgage, most of the monthly payment is applied to interest accruing on the large loan balance. However, over the life of the mortgage, the fraction of the payment covering interest decreases and the fraction reducing principal increases as shown in Figure 3.11. This is a feature common to all amortizing loans including student and auto loans.

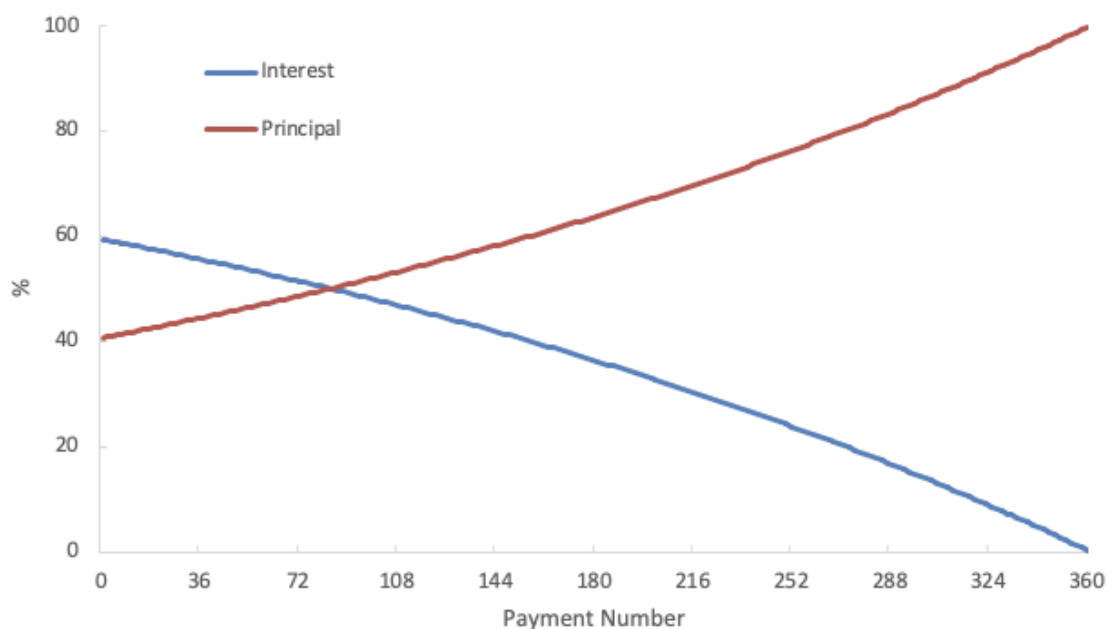


Figure 3.11: Fractions of Monthly Mortgage Payments Going Towards Interest and Principal Reduction

3.5.2 Shopping for a Mortgage

Now that we understand how a mortgage works, let's consider how to choose one. From a financial perspective, there are good and bad mortgages. How do we differentiate? We begin by getting a lot - at least 10 - quotes from different banks, credit unions, and mortgage brokers (people who sell mortgages for multiple financial institutions). The process is simple, requiring nothing more than a phone call or email, and should be done *before* you find a home to buy. The bank will usually want to know your credit status (e.g., strong, average, weak), as well as the geographic area in which you are looking. They'll also need to know

the approximate price of the home and how much you'll need to borrow.⁸

Figure 3.12 presents mortgage quotes on March 19, 2021 for 30-year fixed rate mortgages and a borrower with a good credit history represented by a 730 credit score.⁹ The horizontal axis presents the upfront costs, or points, on the loan. The vertical axis presents the loan interest rate expressed as a monthly compounded APR excluding fees. As borrowers, we like low interest rates and low fees so moving southwest in the graph is preferable. Moving northeast is undesirable.

The red circles correspond to bad loans because for each loan, there is another loan with a lower interest rate and, in some cases, lower fees. The yellow circles, while better than the red circles, are also bad loans because for the same interest rate, we can find a loan with lower upfront costs - further left in the graph. The green circles correspond to good loans because for each of these loans there are no other loans with (i) a lower interest rate and similar or lower costs, and (ii) no other loans with lower costs and a similar or lower interest rate. Put differently, the green circles are the most southwest in the figure. If we drew a curve connecting the green dots (and star), we would find that there are no loans left of or below that curve.

How good are the good loans relative to the bad loans? Let's quantify the difference and in doing so show just how important it is to shop for a mortgage. Take the loan represented by the upper left-most red circle, call it "Red loan," and the loan represented by the lower right-most green circle, call it "Green loan." The Red loan has no upfront fee and an APR 3.5%. The Green loan has a \$14,000 upfront fee and an APR of 2.5%. To make things concrete, we'll assume both loans have a starting principal of \$500,000, we remain in the home for the full 30-year term of the loans, and our opportunity cost is 2% per annum.

⁸What the bank or broker does *not* need to give you a quote is your credit report or credit score. Pulling your credit report or score from one of the credit bureaus - Experian, TransUnion, and Equifax - adversely affects your credit score, though only temporarily. Eventually, whichever lender you choose to go with will need to get that information, but there is no need for it when gathering mortgage quotes.

⁹Credit scores range from 300 to 850. Equifax, one of the three large credit bureaus, describes consumer creditworthiness as falling into five categories.

- 800 to 850: Excellent
- 740 to 799: Very Good
- 670 to 739: Good
- 580 to 669: Fair
- 300 to 579: Poor

The first two, sometimes three, categories are considered **prime borrowers** because of the relatively low credit risk. The last two categories are considered **subprime borrowers**.

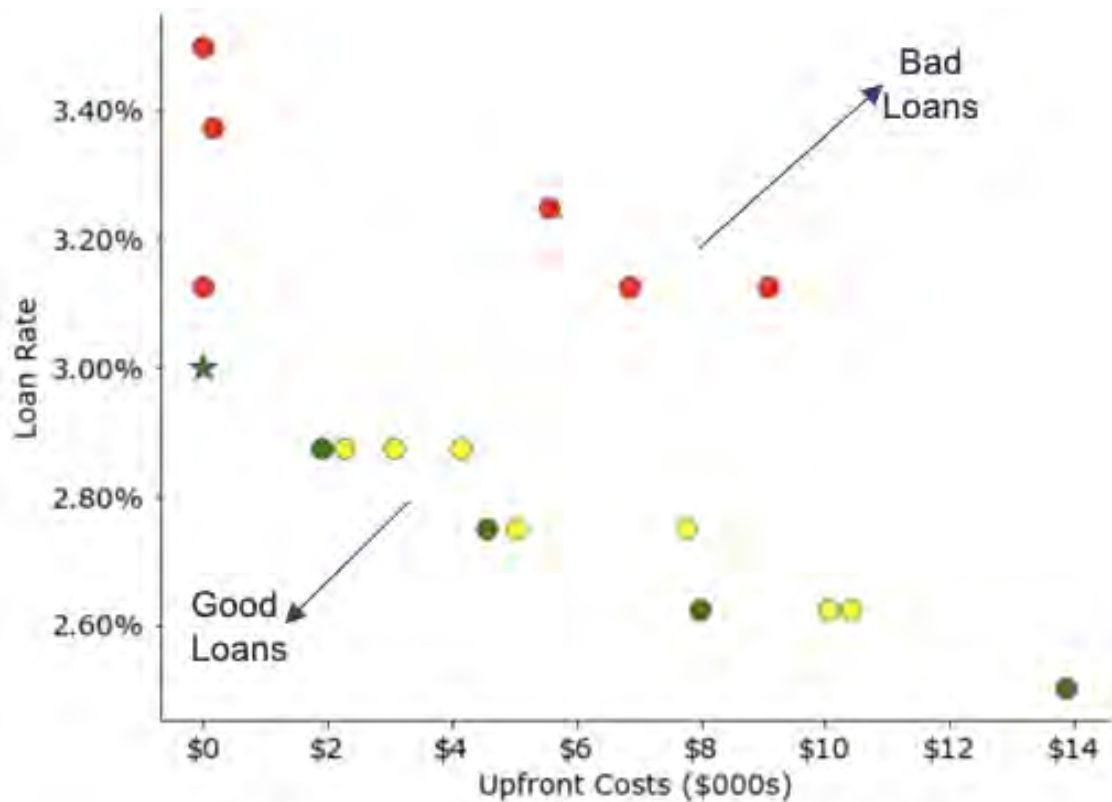


Figure 3.12: Rates and Fees for 30-Year Fixed Rate Mortgages (Source: Bankrate.com, 3/19/2021)

For each loan, let's calculate the net present value (NPV), which tells us how much value each loan creates - positive NPV - or destroys - negative NPV. To do so, we first have to compute the monthly loan payments. Using our annuity cash flow formula (equation 2.2), we get

$$\begin{aligned} \text{Red Loan Monthly Payment} &= \frac{500,000 \times 0.035/12}{1 - (1 + 0.035/12)^{-360}} = \$2,245.22, \text{ and} \\ \text{Green Loan Monthly Payment} &= \frac{500,000 \times 0.025/12}{1 - (1 + 0.025/12)^{-360}} = \$1,975.60. \end{aligned}$$

A useful interpretation of NPV is as the present value (PV) of the benefits minus the present value of the costs. The benefits of the loans is the money we receive today, the

\$500,000 loan principal. The costs are the loan payments and any fees.

$$\begin{aligned}
 PV(\text{Cost of Red Loan}) &= \frac{2,245.22}{(1 + 0.02/12)} + \frac{2,245.22}{(1 + 0.02/12)^2} + \dots + \frac{2,245.22}{(1 + 0.02/12)^{360}} \\
 &= \$607,441.87 \\
 PV(\text{Cost of Green Loan}) &= 14,000 + \frac{1,975.60}{(1 + 0.02/12)} + \frac{1,975.60}{(1 + 0.02/12)^2} + \dots + \frac{1,975.60}{(1 + 0.02/12)^{360}} \\
 &= \$548,496.87
 \end{aligned}$$

The NPV of the Red loan is $500,000 - 607,442 = -\$107,442$. The NPV of the Green loan is $500,000 - 548,597 = -\$48,697$. Both loans are negative NPV and destroy value. That is, the loans are costly for us the borrower. We saw this result above. But, the Red loan is much costlier, $-107,442 - (-48,697) = -\$58,945$. This difference is more than 11% of the total loan principal and a big number in absolute terms. The message of this exercise is simple: Shopping for a mortgage can have a huge impact on the cost of buying a home and our financial well-being.

After eliminating the red and yellow loans, the remaining challenge is deciding among the green loans. The key factor in deciding among these loans is how long we'll be in the house.¹⁰ The green loans present a trade off. Pay more today in upfront fees to save more later with lower monthly payments. The longer we're in the house, the greater the benefit of lower monthly payments and the more valuable the lower interest rate.

The technical appendix provides details on computing the optimal time in the home to select the best mortgage. Figure 3.13 illustrates the intuition. The figure shows the present values of loan costs of remaining in the home for different lengths of time for two different loans. Loan A has a 3% APR and no upfront fees. Loan B has 2.5% APR and \$14,000 upfront fee. There are several aspects of the figure worth noting.

First, both lines are upward sloping meaning the longer we stay in the home, the more expensive both loans become as we continue to pay interest at a rate greater than our opportunity cost (2%). Second, the blue line is steeper than the red line implying the costs for loan A increase more rapidly than those for loan B because of the higher interest rate - 3% vs. 2.5%. Third, the point at which the two curves intersect, 76 months, is the point at which we are indifferent between the two loans because they both have the same cost. Finally, If we plan on staying in the home longer than 76 months, then we should select loan

¹⁰Another consideration is whether we have enough money to pay for the down payment and upfront loan cost. However, we often have the option to "roll the upfront fees" into the loan principal, though this is rarely preferable because our opportunity cost will almost always be less than the the loan interest rate.

B, which has a lower cost (the red line is below the blue line). If we plan on staying in the home less than 76 months, then we should select loan A for the same reason. The intuition is simple. If we're going to pay \$14,000 to reduce the interest rate, we have to stay in the home long enough - at least 76 months - to enjoy the benefit of the lower monthly payments.

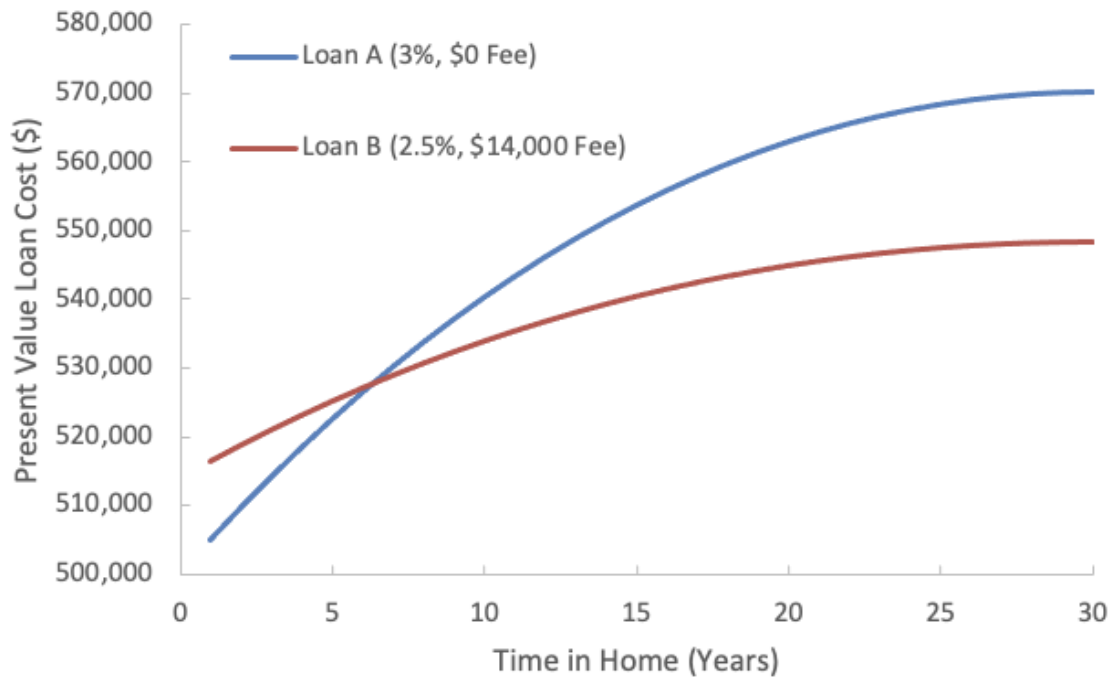


Figure 3.13: Present Value Cost Differential for Two \$500,000 Loans as a Function of Time in the House)

In sum, shopping for a mortgage is important. We can save a great deal of money by selecting the lowest cost mortgage, which depends in part on how long we plan on staying in the home.

3.5.3 Refinancing a Mortgage

To refinance a mortgage means to replace an existing mortgage with a new one. There are several reasons we should consider refinancing. First, refinancing allows us to take advantage of lower interest rates, which can reduce the cost of the mortgage. Second, refinancing allows us to modify the loan term or maturity. Finally, cash-out refinancing allows us to extract equity trapped in our house. In other words, if home prices have risen since we purchased our home, we may be able to borrow even more than what we currently owe on our mortgage and keep the difference, i.e., cash-out. Regardless of the reason, refinancing can have a significant effect on our monthly loan payments and our financial well-being.

However, refinancing is rarely free. Refinancing can come with fees, such as documentation, appraisal, application, title search, and insurance. Some mortgages may have pre-payment penalties so when it comes time to refinance, using the proceeds of the new loan to repay the old one will incur the pre-payment penalty. These costs create a tradeoff with the benefit of lower payments when deciding whether or not to refinance. This should sound familiar. Refinancing a mortgage is really no different than shopping for a mortgage. We're choosing between our existing mortgage and a new one. As such, we'll approach the problem in the same manner by comparing the costs of the two options.

Let's consider our previous example in which we had a 30-year, fixed-rate mortgage with a 3.00% APR and starting balance of \$500,000. Imagine now that five years have gone by so "today," period 0, is just after we've made our 60th mortgage payment. Let's also assume that the interest rate on a new 30-year mortgage is 2.75%, 25 basis points lower than our existing mortgage, but it will cost \$2,500 in upfront fees to get the new mortgage. (One **basis point** is 0.0001 or 0.01%; one **point** is 0.01 or 1%.) Should we refinance our existing mortgage?

Cost of the Existing Mortgage

The value of the existing mortgage is the sum of the remaining mortgage payments discounted by our opportunity cost of capital.

$$\frac{2,108.02}{(1 + 0.02/12)} + \frac{2,108.02}{(1 + 0.02/12)^2} + \dots + \frac{2,108.02}{(1 + 0.02/12)^{300}} = \$497,345.43$$

which can also be computed using our annuity present value formula.

$$\frac{2,108.02}{0.02/12} (1 - (1 + 0.02/12)^{-300}) = \$497,345.43$$

Because these cash flows correspond to payment we have to make, this value is the cost of the mortgage to us.

Value of the New Mortgage

To value the new mortgage, we need first find the new mortgage payments which require knowledge of the outstanding balance on our existing loan. Remember, in a refinancing, we're going to take out a new loan to pay off the old loan. Recall from above that the outstanding balance on a loan is equal to the sum of the remaining payments discounted at

the *loan interest rate*. In our case here, we have 300 remaining payments of \$2,108.02, the value of which is

$$\frac{2,108.02}{(1 + 0.03/12)} + \frac{2,108.02}{(1 + 0.03/12)^2} + \dots + \frac{2,108.02}{(1 + 0.03/12)^{300}} = \$444,531.82.$$

Again, we could also use our annuity present value formula to get the same result.

$$\frac{2,108.02}{0.03/12} (1 - (1 + 0.03/12)^{-300}) = \$444,531.82$$

The monthly payment is computed using the annuity cash flow equation 2.2.

$$\frac{444,531.82 \times 0.0275/12}{(1 - (1 + 0.0275/12)^{-360})} = \$1,814.76$$

Discounting these payments by our opportunity cost gets us the present value of the new loan payments.

$$\frac{1,814.76}{(1 + 0.02/12)} + \frac{1,814.76}{(1 + 0.02/12)^2} + \dots + \frac{1,814.76}{(1 + 0.02/12)^{360}} = \$490,981.15$$

Add to this amount the up front cost of \$2,500 gives us the total cost of the new mortgage, \$493,481.15

The Decision

The existing mortgage will cost us \$497,345 in today's dollars. The new mortgage will cost us \$493,481. Refinancing saves us the difference between these costs, \$3,864, *if we stay in the house for the next 30 years*. Refinancing a mortgage represents a tradeoff. We're paying upfront fees today to lower our monthly payment in the future. We have to stay in the house long enough for the benefits of the lower payment to offset the upfront cost. Exactly how much longer is discussed in detail in the technical appendix.¹¹

3.5.4 Paying Down Your Mortgage Early

Should we use extra money to pay down our mortgage more quickly? The popular press and much of social media would have us believe that paying our mortgage quickly is unequivocally

¹¹Absent from our discussion is the option to wait to refinance our mortgage in the hope of an even lower interest rate next period (month, quarter, year). To properly assess this option requires tools beyond the scope of this text but nicely covered in Pietro Veronesi's text, "Fixed Income Securities: Valuation, Risk, and Risk Management."

better than saving our money. In fact, we're already well equipped to answer the question. If the cost of our loan is greater than our opportunity cost, we should pay down the loan. Otherwise, we should not pay down our loan.¹² Let's use this knowledge to analyze some common arguments in favor of paying down our mortgage early, and then explore the problem more deeply. To do so, we'll continue with our example of a fixed-rate, 30-year mortgage for \$500,000 and with a 3.0% APR.

Common Arguments in Favor of Paying Down Our Mortgage Early

1. We'll be debt free earlier. It depends. Extra payments on our mortgage reduce the loan principal. When we do this, we typically have a choice. We can ask our lender to re-amortize the loan, in which case they will recompute the monthly payment over the same loan term. The result will be a lower monthly payment because we have reduced the principal more quickly with the extra payments. In other words, it takes us just as long to pay off the loan, but our required monthly payments will be smaller. Alternatively, we can ask our lender *not* to re-amortize the loan, in which case we will pay down our mortgage in less time than the original term.

Let's assume that we use our extra money to pay down our mortgage, which is not re-amortized. In other words, we pay $2,108 + 250 = \$2,358$ every month. To figure out how long it will take to pay off our mortgage, we can use equation 2.4 with the periodic interest rate (i) in place of r , and T corresponding to the number of months. Plugging in the values for our mortgage yields

$$T = -\frac{\ln\left(1 - \frac{\text{Value}_t \times i}{CF}\right)}{\ln(1 + i)} = -\frac{\ln\left(1 - \frac{500,000 \times 0.03/12}{2,358.02}\right)}{\ln(1 + 0.03/12)} = 302.48 \text{ months.}$$

By paying an extra \$250 each month, we can pay off our mortgage in a little over 25 years, as opposed to 30 years.

The fractional month, 0.48, means our last payment is a partial payment, which can be calculated as follows. First, estimate the present value of making 302 payments of \$2,358.02 using either our Fundamental Value Relation

$$\frac{2,358.02}{(1 + 0.0025)} + \frac{2,358.02}{(1 + 0.0025)^2} + \dots + \frac{2,358.02}{(1 + 0.0025)^{302}} = \$499,472.39$$

or our present value of an annuity relation

$$\frac{2,358.02}{0.0025} \times (1 - (1 + 0.0025)^{-302}) = \$499,472.39$$

¹²To be precise, there are other considerations, such as the diversity of our investments and wealth, that are relevant and that we'll touch on later in the book.

This estimate shows that 302 payments of \$2,358.02 is $500,000 - 499,472.39 = \$527.62$ short of paying off the loan completely. The future value of this shortfall 303 months from today is

$$527.62 \times (1 + 0.0025)^{303} = \$1,124.31.$$

This amount is the last payment occurring in month 303 to fully pay off the loan.

2. We'll pay less interest. True, but this is not a compelling argument because it ignores the opportunity cost of paying down the mortgage. Additionally, this argument is typically presented in a completely nonsensical manner like so. If we don't pay down our mortgage early, we will pay

$$360 \times 2,108.02 - 500,000 = \$258,887.20$$

in interest. If we do pay down our mortgage early - by paying \$250 extra per month - we only pay

$$302 \times 2,358.02 + 1,124.31 - 500,000 = \$213,246.35$$

in interest. Thus, we'll save $258,887.20 - 213,246.35 = \$45,640.85$ in interest.

Can you see what's wrong with these calculations? They violate rule #1 of finance: Don't add cash flows at different points in time! They have different time units. The figures above are meaningless, yet these calculations are often found in truth in lending and mortgage disclosure forms that report the "total interest" as the sum of all the interest expense. This simple addition assumes that a dollar of interest expense 30 years from today is just as costly to us as a dollar of interest expense one month from today. That's absurd.

3. A psychological burden is lifted earlier. Debt creates worry and stress for most people. This is a fact. Alleviating debt earlier will alleviate the accompanying stress earlier. While this is not a financial consideration, it is nonetheless something that can't be ignored. However, by understanding the tradeoffs involved in this decision, one can alleviate this stress.

Financial Analysis

As noted above, whether or not divert extra money towards paying of our mortgage depends largely on whether the loan interest rate is greater or less than our opportunity cost. In other words, if a dollar of loan principal costs us more than what we can earn on a dollar

investment of *similar risk*, then we should pay down our mortgage. Otherwise, we should invest the extra money.

Assume our opportunity cost of capital is a monthly compounded 5% per year and consider two savings strategies.

1. Save our extra money

The first strategy directs our extra income, \$250 per month, into savings for the entire loan term. Doing so means that in 30 years we will have \$208,064.66 in our savings account, which is computed as follows.

$$\begin{aligned} Value_{360} &= \underbrace{\left(\frac{250}{(1 + 0.05/12)} + \frac{250}{(1 + 0.05/12)^2} + \dots + \frac{250}{(1 + 0.05/12)^{360}} \right)}_{Value_0} \times (1 + 0.05/12)^{360} \\ &= \$208,064.66 \end{aligned}$$

Of course, we could have recognized our savings as an annuity and used our annuity result (equation 2.1) to get the present value of the savings. Alternatively, we could have compounded each cash flow to period 360 and then added. Both approaches will generate the same answer we got above.

2. Use extra money to pay down mortgage

The second strategy uses the extra money to pay down the mortgage more quickly. In this case, we can't start saving until after the mortgage is completely repaid, which occurs 303 months in the future. Figure 3.14 shows the cash flows we can direct to savings in this scenario.

Months	0	1	302	303	304	305	360
Mortgage					2,108	2,108	2,108
Savings			250	250	250	250	
Partial pmt			984				
Total			1,234	2,358	2,358	2,358	

Figure 3.14: Paying of a Mortgage Early Timeline

In month 303, we have a partial payment of \$1,124.31. (See computation above.) This means we'll have $2,108.02 - 1,124.31 = \$983.71$ that we can save plus the \$250 that no longer needs to go towards paying down the mortgage. Each month thereafter, our mortgage is fully paid off. So, we can save both the mortgage amount, \$2,108.02, and the extra cash,

\$250. The future value of these savings in period 360 at our 5% opportunity cost of capital is

$$\begin{aligned} Value_{360} &= \underbrace{\left[1,233.71 + \frac{2,358.02}{0.05/12} \times (1 - (1 + 0.05)^{-(360-303)}) \right]}_{Value_{303}} \times (1 + 0.05/12)^{(360-303)} \\ &= \$152,919.95. \end{aligned}$$

To summarize, the value 30 years from today of our savings when we always save the extra money is \$208,064.66. The value 30 years from today of our savings when we pay down our mortgage with extra money is \$152,919.95. So, saving our money instead of using it to pay off the mortgage more quickly will leave us with $208,064.66 - 152,919.95 = \$55,144.71$ more money in 30 years, or $55,144.71 \div (1 + 0.05/12)^{360} = \$12,342.85$, in today's dollars. Simply put, it's financially smarter to save our money than pay down our mortgage more quickly.

Comments

The comparison above shows that paying off a mortgage early reflects a tradeoff between saving now versus later. If our opportunity cost is greater than our mortgage cost, as it is in this example, then saving now - not paying off the mortgage more quickly - creates value for us. Otherwise, it destroys value for us. As we've said, the interest rate at which people borrow is typically higher than their opportunity cost. This observation begs the question: Does it ever make sense to not pay down our mortgage as quickly as possible? There are two scenarios in which it does.

First, if interest rates rise after we have taken out a fixed-rate mortgage, then the return on a risk-free investment could be greater than the interest rate on our mortgage. This is exactly what occurred in 2022. Rising inflation was met with large increases in interest rates on very safe investments such as bank savings accounts, certificates of deposits (CDs), and Treasury securities.¹³ Many homeowners that had originated mortgages before the interest rate increases were able to take advantage of their low fixed rate mortgage by saving extra money in these safe, high interest rate investments. (One person I know had a 2.75% 30-year fixed rate mortgage at the time the yield on 30-year Treasury bonds was over 4.50%.) In this situation, paying off a mortgage early can be very costly assuming the homeowner can still pay for other necessities like utilities, food, etc.

¹³CDs are savings products offered by most banks and credit unions that require the saver to deposit money for a fixed amount of time, typical between six-months and five years, in exchange for interest.

Second, if we, the homeowners, are willing to take some risk, not paying down the mortgage more quickly makes sense. Consider the following question. If the stock market returns 5.5% per year on average over the next 30 years, half what it has historically returned, but experiences the same volatility (i.e., year-to-year swings), what is the probability that saving extra money is a bad strategy relative to paying down the mortgage more quickly, and how much money do we stand to lose if we do lose money?

Using stock market data over the last 100 years, the probability that saving extra money turns out to be a bad decision is only 26%, if our mortgage rate is 3%. And, when it does turn out to be a bad decision, we will lose an average of \$7,000, 30 years from today. That is, the value of our savings 30 years from today will be \$7,000 smaller than what it would have been had we diverted the extra money towards paying down the mortgage. The other 74% of the time in which saving extra money is a good decision will lead to an average increase in our future savings of \$24,000. Based on this analysis, saving the extra money is not risk-free but still a favorable bet, *as long as the future bears some resemblance to the past*.

Of course, the stock market isn't the only investment option. For example, corporate bonds, especially riskier corporate bonds, often offer expected returns larger than the cost of a residential mortgage. The only question is whether we are willing to accept the risk of our return being lower than the cost of the mortgage. Because most people that have a mortgage also have savings, they are implicitly making a bet that the return on those savings will exceed the cost of their mortgage. Historically speaking, this has more often than not turned out to be a good bet.

3.6 Should You Pay Cash for a Car?

In 1985, the Los Angeles Times published an article entitled, "Should you pay cash for a new car?" The article told the story of Professor Geoffrey Keppel who had saved \$8,239.05 to purchase a car. The car dealer argued that Professor Keppel would be better off investing his savings in a certificate of deposit that pays 8% and borrowing the money at an interest rate of 14.2% to purchase the car.

Suspicious, Professor Keppel decided to compare the total interest he would pay on the loan to the total interest he would earn from investing in the CD. The loan required 48 monthly payments of \$225.97 for a total of \$10,846.56. From this number, he subtracted the principal of the loan, \$8,239.05, to get a total loan interest expense of \$2,607.59. To find the total interest he would earn on the CD, he first calculated the CD balance at the end

of the four years, $8,239.05 \times (1 + 0.08/12)^{48} = \$11,334.18$. From this number, he subtracted the initial investment, \$8,239.05, to get his total interest earnings of \$3,095.13. Professor Keppel concluded he should save his money in the CD and borrow money to buy the car because the total interest earned on the CD, \$3,095.13, was greater than the total interest paid on the loan, \$2,607.59.

Professor Keppel’s intuition for his finding was as follows. “[T]he 14 percent is applied to a declining balance and the 8 percent is to an increasing balance.” Put differently, “over four years, the average outstanding balance of the loan — the average amount on which he’d be paying interest— was only about half the total amount borrowed. His investment, on the other hand, would earn interest on his full deposited principal, plus continually compounded interest throughout the term.”

Frank Sperling, Vice President at Security National Pacific Bank, was quoted as saying: “An easy rule of thumb is if you can earn an interest rate equivalent to half the interest rate on your loan, you’ll come out ahead.” Referred to as Sperling’s rule, the article concludes that “the same analysis, including the same assumptions could probably be applied to any consumer loan, with Sperling’s Rule a good guide to potential consumer advantage.” That is, consumers should always borrow money as long as they can earn a return that is at least half the interest rate on the loan.

Some comments before digging into the claims of this article. First, this is a true story; you can search for the original article. Second, the logic and conclusions are incorrect; you can search for the follow-up article correcting the original. Third, financial authorities should be viewed with a degree of skepticism until the financial logic is laid bare and fully understood. Before searching for these articles, let’s use what we’ve learned to properly analyze Keppel’s decision and understand where Keppel, Sperling, and the LA Times went wrong.

3.6.1 The Correct Way to Approach the Problem

Keppel has two options. Option 1 is purchase the car with cash, which costs \$8,239.05 today. Option 2 is finance the car by borrowing, which costs \$225.97 a month for 48 months.¹⁴ The cash flows for these options are visualized on the timeline in figure 3.15.

¹⁴The monthly car payments are computed using equation 2.2.

$$CF = \frac{Value_0 \times i}{1 - (1 + i)^{-T}} = \frac{8,239.05 \times 0.142/12}{1 - (1 + 0.142/12)^{-(4 \times 12)}} = \$225.97$$




Figure 3.15: Monthly Compounding - Periodic Interest Rate

To compare the costs of these options, we need to sum the present values of all the loan payments to get the total cost of option 2 in today's dollars. The appropriate discount rate for the loan payments is Keppel's opportunity cost, what he can earn on an investment of similar risk such as his safe CD investment, *not* the loan interest rate. Discounting these cash flows by the loan interest and summing will simply get us the principal of the loan, \$8,239.05. Remember, to value cash flows we use the opportunity cost of capital. To get the outstanding balance on a loan, we use the loan interest rate.

Because the loan payments are an annuity, we can use equation 2.1 to get the present value of the loan payments.

$$Value_0 = \frac{CF}{i} \times (1 - (1 + i)^{-T}) = \frac{225.97}{0.08/12} \times (1 - (1 + 0.08/12)^{-(4 \times 12)}) = \$9,256.23$$

Financing the car will cost Keppel \$9,256.23 in today's dollars. Because the loan gives him \$8,239.05 today, the NPV of the loan from Keppel's perspective is -\$1,017.18. This negative NPV means the loan is value-destructive and should not be undertaken. Professor Keppel should pay cash for the car and avoid the high interest rate - relative to his opportunity cost - loan. More simply, no one should borrow money at 14.2% to earn 8%!

3.6.2 Where Did Keppel et al. Go Wrong?

Table 3 presents two tables side by side. The left side presents the monthly interest and account balance if Keppel invests his money in the CD. The right side presents the loan amortization table. The interest earned each month on the CD is computed by multiplying the monthly periodic rate, $0.08/12 = 0.67\%$, by the starting balance. The ending balance is the sum of the starting balance and the interest earned in the month. The loan amortization table works exactly as it does for a mortgage. (See table 2.)

Table 3 shows that more interest appears to be earned on the CD than what is paid on the loan, $3,095.13 - 2,607.59 = \$487.54$. The table also shows that the average balance in the CD is approximately twice that of the average balance of the loan, $9,706.20 / 4,497.15 = 2.16$. This ratio motivates Sperling's logic. Because the average balance in the savings

Certificate of Deposit Account				Loan Amortization Table			
Month	Start Balance	Interest	End Balance	Start Balance	Interest	Payment	End Balance
0			8,239.05				8239.05
1	8,239.05	54.93	8,293.98	8,239.05	97.50	225.97	8,110.57
2	8,293.98	55.29	8,349.27	8,110.57	95.98	225.97	7,980.58
3	8,349.27	55.66	8,404.93	7,980.58	94.44	225.97	7,849.04
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
46	11,110.49	74.07	11,184.56	662.18	7.84	225.97	444.05
47	11,184.56	74.56	11,259.12	444.05	5.25	225.97	223.33
48	11,259.12	75.06	11,334.18	223.33	2.64	225.97	0.00
Total		3,095.13			2,607.59		
Average			9,706.20				4,497.14

Table 3: CD Account and Loan Amortization Tables

account is twice as large as the average balance owed on the loan, Sperling reasons that the interest rate on the savings account only needs to be greater than half that of the loan interest rate for borrowing money to make sense.

This analysis makes several mistakes beginning with the first rule of Fight Club, I mean Finance: Don't add cash flows arriving at different points in time. They have different time units. Computation of the total interest earned and paid and average balances make this mistake. When we compare the present value of paying cash for the car, \$8,239.05, to the present value of financing the car, \$9,256.23, the result is obvious. Financing is a bad idea.

Another way to see the problem is to see if Professor Keppel's savings will generate enough income to pay off his car loan. So, let's assume that every month, Professor Keppel withdraws \$225.97 from his savings in the CD account to make his monthly loan payment. Let's ignore any penalties associated with early withdrawals from the CD to keep things simple. Table 4 details what happens when he follows this strategy.

The CD starts out with the \$8,239.05. In the first month, this money earns $0.67\% \times 8,239.05.67 = \54.93 in interest. At the end of the month, Keppel withdraws \$225.97 to make his loan payment. His balance at the end of the month is just the starting balance plus the interest minus the withdrawal for the loan payment, $8,239.05 + 54.93 - 225.97 = \$8,068.01$. This process continues every month until month 42 when Keppel doesn't have enough money in the CD account to make his current (or future) loan payments. Uh oh...

In fact, at the end of 48 months, Keppel is going to be \$1,399.30 short on his loan payments. The present value of this amount, $1,399.30/(1+0.08/12)^{48} = \$1,017.18$, is exactly

Month	Start Balance	Interest	Loan Payment	End Balance
0				8,239.05
1	8,239.05	54.93	225.97	8,068.01
2	8,068.01	53.79	225.97	7,895.82
3	7,895.82	52.64	225.97	7,722.49
⋮	⋮	⋮		⋮
40	649.50	4.33	225.97	427.86
41	427.86	2.85	225.97	204.74
42	204.74	1.36	225.97	(19.87)
43	(19.87)	(0.13)	225.97	(245.97)
44	(245.97)	(1.64)	225.97	(473.58)
45	(473.58)	(3.16)	225.97	(702.71)
46	(702.71)	(4.68)	225.97	(933.37)
47	(933.37)	(6.22)	225.97	(1,165.56)
48	(1,165.56)	(7.77)	225.97	(1,399.30)

Table 4: Paying for the Loan from the CD Account

what we computed earlier as the difference in the present values between the two options. So, done properly, this analysis leads to the same conclusion. Financing the car is a bad idea because the interest rate on the loan is greater than Keppel's opportunity cost, the interest rate on the CD.

3.7 Leasing a Car

Now imagine we want to buy a car but unlike Professor Keppel, we don't have enough money to purchase it outright. So, we have to borrow money to buy the car and, generally speaking, we have two options: Take a loan and purchase the car, or take a lease and borrow the car. A lease is a loan permitting the lessee to use the lessor's asset, such as a car, airplane, or warehouse. At the end of the lease, the lessee has to return the asset to the lessor, or purchase the asset outright. If you're confused by the language, you're not alone. What's happening is more simply described with an example.

Imagine we want to lease a Porsche Taycan 4S, which as of early 2022 retailed for \$125,600. We could borrow money from a bank to buy the car. But, maybe we don't want to own the car. Maybe we just want to use it for a few years, and we don't want the headache of having to resell it later. Or, maybe we're not sure we want to own the car for a long time

so we'd like to try it out for a few years before deciding whether or not to buy it. What we can do is borrow the car from the dealership, or someone else that owns the car, in exchange for periodic payments. This is a lease. We are the lessee borrowing the car, the dealership is the lessor lending the car. Leases range in duration from several months to several years. At the end of the lease, we return the car or, in some cases, we have the option to purchase it. The key difference between a lease and a loan is that in a lease we don't own the asset. We're just borrowing it. At the end of the lease, we have to return the car or, If we want to keep it, pay for it then.

Our 2022 Porsche Taycan came with the following lease details.

- The lease term (i.e., length) is three years.
- The monthly lease payments are \$1,599.32 payable at the start of each month.
- Our upfront costs are a \$1,095 acquisition fee, a \$11,885 capital cost reduction payment, and the first lease payment.
- At the end of the lease three years from today, we can purchase the car for its **residual value**, which is \$65,312

If instead we choose to buy the car, we can take out a three year loan at a 4% APR, which requires that we put down 20% of the purchase price of the car. In other words, we would pay $20\% \times 125,600 = \$25,120$ today at signing. This means we need to borrow $125,600 - 25,120 = \$100,480$. At a 4% APR, this implies monthly payments of

$$\frac{100,480 \times \frac{0.04}{12}}{1 - \left(1 + \frac{0.04}{12}\right)^{-36}} = \$2,966.57.$$

This calculation is just our annuity cash flow result (equation 2.2).

3.7.1 Lease vs. Buy and Borrow

Let's assume that our choice of financing - lease or borrow and buy - has no impact on how we drive the car or its future value.¹⁵ What we want to understand is how the two options compare. To do so, let's start by laying out the cash flows to each option. Figures 3.16 and 3.17 display the timelines for each. Focusing on the cash flows at the bottom of each figure, we can see that both correspond to borrowing money. The cash flow today is positive and the future cash flows are negative.

Months	0	1	2	...	35	36
Car price	125.6					
Acquisition fee	-1.1					
Cost reduction	-11.9					
Lease payments	-1.6	-1.6	-1.6		-1.6	
Disposition fee						-0.6
Residual value						-65.3
Cash flows	111.0	-1.6	-1.6		-1.6	-65.9

Figure 3.16: Porsche Taycan Lease Cash Flows (\$000s)

Months	0	1	2	...	35	36
Car price	125.6					
Down payment	-25.1					
Loan payments		-3.0	-3.0		-3.0	-3.0
Cash flows	100.5	-3.0	-3.0		-3.0	-3.0

Figure 3.17: Porsche Taycan Buy and Borrow Cash Flows (\$000s)

When we purchase the car, we are borrowing less today - \$100,480 vs. \$111,021 - but making larger monthly payments - \$1,599 vs. 2,967. However, at the end of the loan term, we own the car whereas with the lease we have to return it or purchase it for its residual value - \$65,312. Thus, the lease payments may be smaller than the loan payments, but there is a really big payment - either the cash to buy the car or the cash we lose by returning it - at the end of the lease.

Let's start by estimating the interest rate on the lease or the one discount rate such that the present value of the future payments equals the amount borrowed.

$$111,020 = \frac{1,599}{1+r} + \frac{1,599}{(1+r)^2} + \dots + \frac{1,599}{(1+r)^{35}} + \frac{65,907}{(1+r)^{36}} \implies r = 0.34\%$$

This is a monthly rate that corresponds to a $12 \times 0.34 = 4.08\%$ APR, slightly higher than the loan APR of 4%. Sometimes leases will quote a **money factor** from which we can estimate the APR by multiplying by 2,400. The money factor for the Porsche lease is $4.08 / 2,400 = 0.0017$.

Now let's compute the net present values of the two options using an annual opportunity

¹⁵We might think that when we lease a car, we don't have to take as good of care of the car because we can return it to the lessor. However, most auto leases limit the annual mileage, charging a penalty if the lessee goes over the limit, and charge for any damages to the car when it is returned.

cost of capital of 3%, or $(1+0.03)^{1/12}-1 = 0.25\%$ on a monthly basis. Again, our opportunity cost is assumed to be lower than the lease and loan rates because in most instances consumers borrow at a higher rate than their opportunity cost.

$$\begin{aligned} \text{NPV Lease: } & 111,021 - \frac{1,599}{1+0.0025} - \dots - \frac{1,599}{(1+0.0025)^{35}} - \frac{65,907}{(1+0.0025)^{36}} = -2,859 \\ \text{NPV Loan: } & 100,480 - \frac{2,967}{1+0.0025} - \dots - \frac{2,967}{(1+0.0025)^{35}} - \frac{2,967}{(1+0.0025)^{36}} = -1,530 \end{aligned}$$

The NPV of the loan is greater than that of the lease suggesting the loan is a financially preferable option.

However, we need to mention what's missing from this analysis. First, we have ignored the value of the option to buy the car at the end of the lease. Valuing this option requires tools beyond the scope of this text, but its value will increase the NPV of leasing. Second, we have ignored the headache of trying to sell our car if we want a new one after three years. We have to take time to meet with potential buyers. We may have to advertise that our car is for sale. We expose ourselves to the risk of not being able to sell our car for some time or at a steep discount if the market for used cars is in a bad state. All of this is to say that the analysis above is more a starting point for the decision to lease vs. buy as opposed to the final word.

3.7.2 Lease Payment Calculations

For those interested, car lease payment typically are not computed using our Fundamental Value Relation. Instead, practitioners use the following formula.

$$\begin{aligned} \text{Lease payment} = & \frac{\overbrace{\text{Adjusted capitalized cost} - \text{Residual value}}^{\text{Depreciation}}}{\text{Lease term}} \\ & + \frac{\overbrace{(\text{Adjusted capitalized cost} + \text{Residual value}) \times \text{Money factor}}^{\text{Finance charge}}}{\text{Lease term}} \\ & + \frac{\overbrace{(\text{Depreciation} + \text{Finance charge}) \times \text{Tax rate}}^{\text{Tax}}}{\text{Lease term}} \end{aligned}$$

The first term estimates how much value the car loses each period, i.e., month. For our Porsche, the adjusted capitalized cost is the price of the car, \$125,600, less our upfront payment, \$11,885, and a federal tax credit of \$7,500 we get because the car is a plug-in electric vehicle. So, the adjusted capital cost is \$106,215. Subtracting the residual value

of \$65,312 and dividing by the 36 month term produces a monthly depreciation charge of \$1,136.19.

The second term estimates the periodic interest or finance charge. The money factor, as discussed above, is the APR times 2,400. Assuming Porsche is targeting a 6.5% interest rate on the lease, then the money factor is $6.5 / 2,400 = 0.0027$. The monthly finance charge is therefore

$$(106,215 + 65,312) \times 0.0027 = \$465.12$$

The final term accounts for any taxes, which we ignored in our Porsche example.

Adding up the depreciation and finance charges produces a monthly lease payment of $1,136.19 + 464.55 = \$1,600.75$, very close to the \$1,599.32. We might wonder why we used a 6.5% APR when we computed a lease APR above of 4.08%. The reason is that our previous calculation accounted for fees, when in practice fees are treated separately from the lease payments. Consider the the least payments as illustrated in figure 3.18, which excludes the acquisition and disposition fees. The implied monthly interest rate for these cash flows is

$$104,616 = \frac{1,599}{1+r} + \frac{1,599}{(1+r)^2} + \dots + \frac{1,599}{(1+r)^{35}} + \frac{65,312}{(1+r)^{36}} \implies r = 0.54\%.$$

The corresponding APR is $12 \times 0.54 = 6.51\%$, very close to the 6.50% we used to get the least payment.

Months	0	1	2	⋮	35	36
Car price	125.6					
Tax credit	-7.5					
Cost reduction	-11.9					
Lease payments	-1.6	-1.6	-1.6	⋮	-1.6	
Residual value						-65.3
Cash flows	111.0	-1.6	-1.6	⋮	-1.6	-65.3

Figure 3.18: Porsche Taycan Lease Cash Flows Ignoring Fees (\$000s)

3.8 Credit Cards

The amortizing loans we've considered thus far have the following similarities. They give us money and we repay them in equal installments over a fixed term. **Credit cards** operate very differently. We can borrow money as often as we like an in any amount we like as long

as our total borrowings do not exceed a certain limit. We can repay as much as we like, as frequently as we like, as long as the frequency of payments is at least monthly and the amount is above a minimum threshold. Credit cards are an example of **revolving credit**, which allows borrowers to draw down and payback money multiple times over the life of the loan.

Unlike other revolving loans, such as those made to corporations, credit cards don't expire. As long as we make the required payments and pay any necessary renewal fees we can continue borrowing and repaying for as long as we like. Despite these differences, using a credit card for purchases is just borrowing money and follows the same fundamental principles we've been discussing all along.

3.8.1 Institutional Details

Like a mortgage or auto loan, credit card debt has to be repaid. So, like a mortgage and an auto loan, similar risk means a risk-free investment like a bank savings account or certificate of deposit. Unlike a mortgage or auto loan, credit card debt is unsecured; there is no asset like a house or car that the lender can seize if we fail to repay the debt. **Defaulting** on (i.e., not paying) a mortgage or auto loan leads to the lender taking our home or car.

Because the credit card is unsecured, the interest rate is higher than that on a mortgage or auto loan - typically a lot higher as seen in Figure 3.1. So, paying off credit card debt, or any high interest rate debt, is often a smart financial strategy, where "high interest" simply means greater than what we can earn on other (near) risk-free investments.

There are a few key things to understand about credit cards beginning with multiple interest rates depending on when you borrow and the reason for the borrowing.

- **Purchase** interest rate applies to all purchases made on the card.
- **Balance Transfer** interest rate applies to all balance transfers from other credit cards.
- **Cash Advance** interest rate applies to all cash withdrawn from a bank or ATM using the card.
- **Penalty** interest rate is applied when you miss a payment and is often higher than the other interest rates (e.g., 29.9%).
- **Introductory** interest rate is often a very low interest rate (e.g., 0%) offered for a limited time (e.g., 10 months). These temporary, low interest rates are sometimes referred to as **teaser rates**.

The introductory rates open up avenues for all sorts of schemes to take advantage of these temporary low rates.

For example, if a card offers an 0% introductory rate on *all* borrowings, we can take a cash advance and save the money in an interest bearing account (e.g., savings, money market, Treasury securities). We make the monthly payments, and just before the introductory rate period ends we can pay the entire card balance. We're left with the interest earned on our savings at no cost to us. Making money without using any of our own money is called an **arbitrage**.

Cash withdrawals often come with fees on top of annual card fees. Additionally, most credit card companies typically only offer the low teaser rates on balance transfers from other credit cards or new purchases. An indirect cost of opening new credit card accounts is the negative impact on our credit score because of the required credit checks, though this is usually temporary.

Ultimately, credit cards are just another type of loan, though one often with very high interest rates. As such, paying off your credit card as quickly as possible is financially often the best course of action. Let's see what can go wrong by investigating the behavior of a former student.

3.8.2 Ryan's Bad Idea

Ryan has a \$5,124.86 balance - all from purchases - on his credit card, which carries an 18.99% APR. The minimum payment due is \$51.00 - which is computed as the larger of 1% of the card balance and \$40. To get his personal finances in order, Ryan is considering shredding his credit card and making the minimum payment each month. When will Ryan be able to pay off his debt?

To answer this, we need only recognize that Ryan is establishing an annuity. He's going to making monthly payments of \$51.00 until all \$5,124.86 - the present value - is paid off. The discount rate is just the monthly periodic rate, $0.1899 \div 12 = 0.015825$. What we don't know is T - how long the annuity will last. However, we can compute T as we did earlier since his payment scheme is just an annuity.

$$\begin{aligned}
T &= -\frac{\ln\left(1 - \frac{Value_t \times i}{CF}\right)}{\ln(1 + i)} \\
&= -\frac{\ln\left(1 - \frac{5,124.86 \times 0.015825}{51.00}\right)}{\ln(1 + 0.015825)} \\
&= \text{\#NUM!}
\end{aligned}$$

What the heck is “#NUM!”?!?!? That’s what will show up in Excel if you execute this computation. The problem is, the formula doesn’t make any sense in this situation. To see why, consider what the interest expense would be at the end of the month on a balance of \$5,124.86.

$$\text{Interest expense} = 5,124.86 \times 0.015825 = \$81.10$$

In other words, the interest expense is larger than the minimum payment. But, this means the balance will continue to grow over time. So, paying \$51.00 every month not only won’t reduce what Ryan owes, he’ll actually wind up owing more and more each month! Look at the first few lines of the amortization table 5

	Start			End
	Balance	Interest	Payment	Balance
0				5,124.86
1	5,124.86	81.10	51.00	5,154.96
2	5,154.96	81.58	51.00	5,185.54
3	5,185.54	82.06	51.00	5,216.60
⋮	⋮	⋮	⋮	⋮

Table 5: Credit Card Amortization Table

We might think: “Great! We’ll just pay \$51 each month and die with the debt.” Unfortunately, credit card companies are on to this scheme. What will happen is that the minimum payment next month will again be the larger of 1% of the balance and \$40 *plus* the interest that accrued over the month (i.e., the \$81.10). So, eventually, he’ll wind up paying down the credit card debt, just over a very long period of time. And, because the effective annual rate he’s paying is $(1 + 0.1899/12)^{12} - 1 = 0.2073$, this strategy is almost surely a terrible idea because it’s very unlikely to find a risk-free investment offering that high of a return.

3.9 Key Ideas

We learned how to use our Fundamental Value Relation when cash flows come and go and frequencies other than once a year. We also learned some new finance jargon.

- Loans are costly to borrowers because the loan interest rate is typically larger than the borrower's opportunity cost. However, when the borrower's opportunity cost is greater than the loan interest rate, borrowing money can be financially beneficial.
- The compounding frequency tells us how often interest is compounded within a year. See table 1 for a list of common frequencies.
- There are three types of interest rates: APR, EAR (a.k.a., APY) denoted r , and periodic interest rate denoted i . These rates are linked by equations 3.1 and 3.2. The EAR is a discount rate when we measure time in years. The periodic interest rate is a discount rate when we measure time in periods (month, quarter, semi-annual, etc.) The APR is a interest rate quote, *not* a discount rate, and as such should not be used to discount cash flows unless the payments and compounding frequency are annual, in which case the APR, EAR, and periodic interest rate are all equal.
- Compound interest refers to interest earning interest and is a consequence of compounding.
- Mortgages, leases, student loans, auto loans, and credit cards are all just loans with monthly interest accrual and monthly payments. The first four are amortizing loans because principal is gradually reduced of the life of the loan due to the fixed payment size. Credit cards are a type of revolving loan in which the loan can be drawn upon and pay back multiple times.

3.10 Technical Appendix

3.10.1 Selecting a Mortgage Based on Time in the Home

Consider two fixed-rate loans, A and B, with the same maturity and monthly payments of CF_A and CF_B , respectively. Loan A has no upfront costs and periodic interest rate i_A . Loan B has upfront costs of k and periodic interest rate i_B . Because of the upfront costs, i_B is less than i_A . The borrower's monthly opportunity cost is i . The present values of the costs of

the two loans if the house is sold in some month t that is less than or equal to the maturity of the loans, T , are:

$$\begin{aligned}
 \text{Present value loan A costs} &= \overbrace{\frac{CF_A}{i} (1 - (1 + i)^{-t})}^{\text{PV of monthly payments}} + \overbrace{\frac{CF_A}{i_A} (1 - (1 + i_A)^{-(T-t)}) \times \frac{1}{(1 + i)^t}}^{\text{PV of principal at } t} \\
 \text{Present value loan B costs} &= \underbrace{k}_{\text{upfront cost}} + \overbrace{\frac{CF_B}{i} (1 - (1 + i)^{-t})}^{\text{PV of monthly payments}} \\
 &\quad + \overbrace{\frac{CF_B}{i_B} (1 - (1 + i_B)^{-(T-t)}) \times \frac{1}{(1 + i)^t}}^{\text{PV of principal at } t}
 \end{aligned}$$

Equating these two expressions and solving for t produces the time in the home, measured in months, for which the loans are equal. For stays less than t , loan A is preferable because it avoids the upfront cost. For stays greater than t , loan B is preferable because it takes advantage of the lower interest rate (and monthly payments). Solving for the t that equates the two costs can be done using Excel's Goal seek or Solver functions.

3.10.2 Continuously Compounded Interest Rates

What happens when interest is **continuously compounded**? In other words, what happens when the compounding period gets arbitrarily small? Think every nanosecond but even quicker. The annual rate becomes

$$\lim_{k \rightarrow \infty} \left(1 + \frac{APR}{k} \right)^k = e^{APR} \quad (3.5)$$

where e is Euler's number approximately equal to 2.71828. The T -year rate of return is

$$\lim_{k \rightarrow \infty} \left(1 + \frac{APR}{k} \right)^{kT} = e^{APR \times T} \quad (3.6)$$

For example, if the APR is 5% and compounding is continuous, then the effective annual rate is $e^{0.05} - 1 = 0.051271$, which is slightly greater than daily compounding $(1 + 0.05/365)^{365} - 1 = 0.051267$ or 5.13%. Over a 10-year horizon, the continuously compounded return is $e^{0.05 \times 10} - 1 = 0.64872$ or 64.87%. Over a one-quarter horizon, the continuously compounded return is $e^{0.05 \times 0.25} - 1 = 0.012578$ or 1.26%. Notice that we measure time in years; one quarter is 0.25 of a year.

While interest is not continuously compounded in real life, assuming so makes some financial calculations easier. We won't use it much in this book, but familiarity with it is useful in other settings.

3.11 Problems

- 3.1 (*Saving to pay a loan*) You are required to pay \$100,000 five years from now to your local bank for a loan. How much money do you need to invest today in an account paying 3% APR in order to have enough money for the \$100,000 payment if interest on your savings is compounded:
- annually,
 - quarterly,
 - monthly, and
 - daily
- 3.2 (*Mortgage payments*) What is the monthly payment on a 30-year fixed-rate mortgage for \$250,000 with an APR of 9% compounded monthly?
- 3.3 (*Interest rates on credit cards*) According to Wallethub, the highest APR on an active US consumer credit card is 36% and offered by First Premiere Bank. What are the corresponding periodic and effective annual rates if interest is compounded monthly?
- 3.4 (*Car loan*) You are preparing to purchase a new McLaren 600LT for \$320,000. Your plan is to pay \$70,000 in cash and borrow the remaining \$250,000 from the dealer, who is offering a 72-month loan with an APR of 4.9%. Loan repayment occurs in equal, monthly installments over the term of the loan.

Using this information answer the following questions:

- What is the monthly periodic interest rate?
 - What is the effective annual interest rate (EAR)?
 - What is the monthly payment?
 - How much interest will you pay in total over the life of the loan (ignore the time value of money)?
- 3.5 (*Bank loan*) Your local bank is offering to lend you \$500,000 for 5-years at a 6.25% APR. With this information, answer the following questions:

- a. What is the quarterly periodic interest rate?
 - b. If the loan must be repaid in equal quarterly installments, how large will the installments (i.e., quarterly payments to the bank) be?
- 3.6 (*Valuing an annuity*) An insurance company is selling an annuity that will pay \$12,000 per month for 20 years. If the current APR is 3.5%, what is the value of this annuity?
- 3.7 (*Saving for a car*) One and a half years from today, you would like to buy a new car that requires a down payment of \$5,000. Using this information, answer the following questions.
- a. How much do you have to save today to have enough for the down payment assuming your bank is offering an APR of 3.2% with quarterly compounding?
 - b. A local competitor bank is offering a special 3.4% APR with annual compounding. Should you switch banks?
- 3.8 (*Mortgage payments*) You are buying a new home and you can only afford monthly payments of \$2,000. How large of a loan can you afford if:
- a. the term is 15 years and the APR is 6.4%,
 - b. the term is 30 years and the APR is 7.9%,
- 3.9 (*Savings account interest rates*) Marcus is the retail banking arm of Goldman Sachs. As of January 17, 2021, Marcus was offering a savings account with an annual percentage yield (APY) of 0.50% and daily compounding. APY is a term used in practice to measure the amount of interest earned in one year accounting for the effects of compounding. Using this information, answer the following questions:
- a. What is the effective annual rate (EAR) on the account?
 - b. What is the annual percentage rate (APR) on the account?
 - c. What is the periodic interest rate on the account?
 - d. If you deposit \$25,000 today and make no subsequent deposits or withdrawals, how much money will be in your account five years from today?
- 3.10 (*Staggered annuity*) Your parents make you the following offer. They will give you a \$500 monthly allowance starting next month for five years - a total of 60 payments. In return, you must repay them \$550 a month for the following ten years beginning one

year after you receive the last payment from your parents. Assume the opportunity cost of capital is 12% per annum with monthly compounding. Should you accept their offer?

3.11 (*Comparing savings account rates*) Which savings rate option do you prefer and why?

- a. 5% APR with annual compounding
- b. 2.5% every six months
- c. 4.75% APR with daily compounding

3.12 (*Implied interest rate*) A local bank is running the following advertisement in the newspaper: “For just \$20,000 we will pay you \$100 forever!” The fine print in the ad says that for a \$20,000 deposit, the bank will pay \$100 every month in perpetuity, starting one month after the deposit is made. What is the interest rate, expressed as an APR, implied by this investment opportunity?

3.13 (*Loan down payment-repayment relations*) You are preparing to purchase a new McLaren 600LT for \$320,000. Your plan is to pay \$70,000 in cash and borrow the remaining \$250,000 from the dealer, who is offering a 72-month loan with an APR of 4.9%. Loan repayment occurs in equal, monthly installments over the term of the loan.

Using this information answer the following questions:

- a. What is the monthly car payment?
- b. Create a plot showing the relation between the down payment (i.e., cash payment) and the monthly payment. Put the down payment on the horizontal axis and the monthly payment on the vertical axis.
- c. (*Challenging*) What is the slope of the relation in the previous question and its interpretation?

3.14 (*Credit card minimum payment and duration to repay*) You currently owe \$12,475 on your credit card which charges interest at an APR of 17.75% with monthly compounding. The monthly minimum payment is computed as 2% of the balance owed. Answer the following questions.

- a. What is the current monthly minimum?
- b. If you continually make the current monthly minimum payment, how long will it take to pay off your credit card?

c. (*Challenging*) To avoid overpaying, how large will your last payment be?

3.15 (*Insurance payment plans*) Scott is considering purchasing a life insurance policy. The insurance company offers several different payment options detailed in the following table.

Payment Frequency	Amount (\$)	Cost per Year (\$)
Annual	2,005.50	2,005.50
Semi-annual	1,032.83	2,065.66
Quarterly	525.44	2,101.76
Monthly	\$175.49	2,105.88

For each plan, the first payment is due immediately at the start of the policy. Assume Scott's annual opportunity cost of capital is 6% and that he has enough money to pay the policy premium in full at the start of the policy.

Using this information, answer the following questions.

- How is the cost per year calculated?
- What are the implied periodic rate, APR, and effective annual rate for each payment plan?
- What is the cost of each plan in today's dollars?
- Which plan should Scott select? Why?
- Assuming Scott's annual opportunity cost of capital is 13%, which plan should he select? Why?

3.16 (*Mortgage mechanics, amortization schedule*) Lena has just purchased her first home. The home cost \$945,000. She paid \$300,000 in cash and the rest with a mortgage. The mortgage is a 30-year, fixed-rate loan with a 3.00% APR (monthly compounding).

With this information, answer the following questions:

- What is the principal amount of the loan?
- What is the monthly mortgage payment?
- How much money is owed on the loan after the 83rd payment?
- What fraction of the 84th payment is interest? Principal reduction?
- (*Challenging*) How long will it take to repay one half of the loan?

- 3.17 (*Credit card repayment requirements*) You have an \$8,000 balance on your credit card, which has an APR of 18% with monthly compounding. How long will it take you to pay off your credit card if you pay \$100 per month? What is the minimum payment the credit card company would require you to make, assuming your debts were passed on to others indefinitely?
- 3.18 (*Credit card mechanics*) D.J. has a current balance of \$2,000 on his credit card, which charges a 24% APR with monthly compounding. He can afford to pay \$100 per month. What will his card balance be one year from today, assuming his first payment occurs one month from today?
- 3.19 (*Investment financing terms*) Strontium Inc is a mining company looking to purchase a new Cat 797F mining truck. The truck retails for \$3.4 million. Caterpillar offers financing that requires \$55,000 monthly payments over a 72-month term. Strontium's cost of capital - i.e., how much it costs to raise money for its investments - is quoted as a 6.8% APR with semi-annual compounding. (Typically, a firm's cost of capital is expressed as an EAR, but let's play along for the sake of skill-building.)

Using this information, answer the following questions.

- a. What is Strontium's semi-annual cost of capital?
 - b. What is Strontium's effective annual cost of capital?
 - c. What is Strontium's monthly cost of capital?
 - d. Should Strontium finance the acquisition of the truck or purchase it outright?
 - e. (*Challenging*) Strontium is looking to push its payments later to conserve liquidity today, but does not want negatively impact the value of the loan. Caterpillar is happy to accommodate as long as the payments grow at a constant rate, g , from an initial payment of \$25,000 at the end of the first month. What growth rate, g , will ensure that Caterpillar's and Strontium's concerns are addressed? How large will the last payment be under this financing arrangement?
- 3.20 (*Annuity valuation*) You are considering purchasing a 10-year annuity with quarterly payments of \$500 beginning one quarter from today. The annuity APR is 3.75% with semi-annual compounding?

Using this information, answer the following questions.

- a. What is the quarterly periodic rate?
- b. What is a fair price for the annuity?

- c. (*Challenging*) If you decide to sell the annuity five years from today after receiving the 60th payment, what would be a fair price for the annuity at that time assuming the APR has not changed?

3.21 (*Retirement savings, dynamic investment strategies*) You have a goal of entering your retirement, 30 years from today, with \$5,000,000. To do so, you are constructing an annual savings strategy in which the first contribution will occur one year from today and the last 30 years from today.

Your investment strategy over the 30 years is as follows:

- Years 1-5: Aggressive equity investment with 11% annual expected return.
- Years 6-20: Blended bond and equity investment with 8% annual expected return.
- Years 21-30: Conservative fixed income investment with 5% annual expected return.

Using this information, answer the following questions.

- a. What is the value today of the \$5,000,000 goal, thirty years from today?
- b. (*Challenging*) What constant amount of money must you save each year to achieve your nest egg goal of \$5 million in thirty years?
- c. (*Challenging*) How much must you save beginning next year if you allow your savings to grow by 3% each year?
- d. (*Challenging*) What are the algebraic relations for your solutions to the previous two questions?

3.22 (*Mortgage mechanics, interest tax shield*) Ten years ago you purchased a \$150,000 investment property with \$15,000 down and a 25-year fixed-rate mortgage with an APR of 8.4% and monthly compounding.

Using this information, answer the following questions.

- a. What is the monthly mortgage payment?
- b. What is the outstanding loan balance today?
- c. (*Challenging*) How much interest is paid in the 10th year - months 108 through 120 - ignoring the time value of money. (A spreadsheet is really helpful here.)
- d. (*Challenging*) What are the tax savings from this interest expense if the relevant tax rate is 25% and we itemize our taxes so that this interest is tax deductible.

3.23 (*Payday lending, implied interest rates*) You are short on funds this month for groceries and utilities. To hold you over until you receive your paycheck, you turn to American Financial, which offers **payday loans**. You decide to borrow \$500, which must be repaid along with a \$25 fee one month from now.

Using this information, answer the following questions.

- a. What is the implied monthly periodic interest rate of the loan?
- b. What is the implied APR?
- c. What is the implied EAR?
- d. At the end of 30 days, you can **roll over** your loan into a new loan. That is, instead of paying the \$525 back at the end of the 30 days, you can extend the loan for another 30 days. If you do, there will be an additional fee of \$45 so the total payment in 60 days will be \$570. What are the implied monthly periodic interest rate, APR, and EAR of this loan?

3.24 (*Savings and inflation*) Derrick Fielding is saving money to purchase a new Audi Q8 - it's a car - that currently costs \$87,500. Auto industry experts believe auto prices will increase with inflation, which is currently projected to be 4.0% per year for the foreseeable future.

Derrick's bank offers two certificates of deposit (CDs) in which he can invest his money. The term or duration of each CD, and corresponding APR are detailed in the table. Interest is compounded monthly.

Term (Years)	APR (%)
1	3.0%
2	5.0%

Using this information, answer the following questions.

- a. What are the expected prices of the Audi one and two years from today?
- b. What are the periodic and effective annual interest rates for the 1- and 2-year CDs?
- c. Compute the real annual returns Derrick will earn on the one-year and two-year CD investments.

- d. Assuming the price of the Audi increases with the rate of inflation, how much money must Derrick invest today if he wishes to purchase the car one year from today and invests in the one-year CD? How would your answer change if he chose to purchase the car two years from today and invests in the two-year CD?

3.25 (*Auto financing*) Robert Michaels is looking to purchase a GMC Yukon Denali, a large SUV for the wife, two kids, two dogs, and gear. The car he's spec'd out costs \$89,500 before taxes and fees. GMC is offering a five-year, fixed-rate auto loan with a 4.69% APR. Interest is compounded monthly to coincide with the monthly loan payments. The first loan payment is due one month after signing. Robert's opportunity cost of capital is 3.5% per year.

Using this information, answer the following questions.

- a. What are the periodic and effective annual interest rates on the loan?
- b. If Robert puts down \$5,000 in cash and borrows the rest of the vehicle cost - ignoring taxes and fees - what is his monthly loan payment?
- c. The dealer has offered to roll (i.e., include) the taxes and fees into the loan to save Robert money at signing. Assuming Robert has the money to pay for the taxes and fees, should he take advantage of this offer and save his money for other uses? Please answer yes or no and provide a brief explanation for your answer.
- d. How much will Robert owe on the loan one year after purchase, just after his 12th payment?
- e. One year after buying the car, Robert plans to use his year-end bonus of \$25,000 to pay down his car loan. What is his new monthly payment, assuming his loan servicer re-amortizes the loan over the original term (i.e., he still has 48 payments to make)?
- f. (*Challenging*) Continuing from the previous question, how long would it take Robert to pay off the loan if instead of re-amortizing the loan, Robert continued making the same monthly payments you computed in problem 2 above?
- g. (*Challenging*) What would the monthly payment be if the first payment was due at signing, immediately, instead of one month from the signing of the loan? (Assume again that Robert has \$5,000 in cash and that the last payment occurs 5 years from today.)

3.26 (*Reverse mortgage*) (*Challenging*) Letty and Jerry Katz are 72 year-old retirees. Their fixed income enables them to pay for all of their non-discretionary expenses (e.g., food, taxes, insurance, health care), but they are unable to afford to travel as much as they would like. As such, they have decided to get a fixed-rate reverse mortgage to tap into their home equity (i.e., to get money out of their home without selling it).

Based on the value of their home (recently appraised at \$1.2 million), their age, and their credit risk, they have been approved for a lump sum payment today of \$672,000 today, which is net of fees and closing costs. The monthly compounded APR on the loan is 6.680% but there are no monthly payments to make. Rather, the loan is repaid with the proceeds from the sale of the home some time in the future. The loan is non-recourse in that if the proceeds from the sale of the home are less than the mortgage, the lender has no recourse to recover the shortfall.

The Katz are expected to remain in their home for 14 years and the price of their home is expected to increase with inflation at 3% per year.

Using this information, answer the following questions.

- a. Assuming just for this question that the price of the Katz's home *doesn't* change, how long can the Katz stay in the home before the lender loses money on the loan (i.e., the value of the home is less than what is owed on the loan)?
- b. How long can the Katz's stay in the home before the lender loses money on the loan (i.e., the value of the home is less than what is owed on the loan) if home price appreciation is as described above? What is the value of the home at this point in time?
- c. What is the lowest annual home price appreciation over this period before the lender begins to lose money, assuming the Katz's remain in the home for 14 years?
- d. What is the highest interest rate the lender can charge on the loan before expecting to lose money, assuming the Katz's remain in the home for 14 years?
- e. What is the lender's expected profit at the expected time of sale, 14 years from now? What is the corresponding annualized expected return on their investment?
- f. The lender has a 10% annual hurdle rate - expressed as an EAR - on their investments, meaning they expect to earn at least 10% per year on their investments. How large of a loan can the lender make to the Katz's and still achieve this hurdle rate?

3.27 (*Private school pricing*) Andreea is considering sending her daughter to Caldwin, a private school located in Connecticut. Tuition is \$48,000 per year to attend, and the school offers three payment options.

-
- a. Pay in full on July 1
 - b. Pay 60% of the tuition on July 1, and the remaining 40% on December 1
 - c. Pay \$4,050 per month starting July 1 for 12 months
-

Andreea's bank offers a savings account paying a daily compounded APR of 5.15%.

Using this information, answer the following questions assuming today is July 1.

- a. What is the implied monthly interest rate for option b.? What is the corresponding APR?
- b. What is the implied monthly interest rate for option c.? What is the corresponding APR?
- c. How much does each option cost Andreea in today's (July 1) dollars? Which option should she select?

Part II

Decisions Most People Make

Chapter 4

A Financial Perspective of Business

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

Now we're going to transition into financial decision making in business settings. Fortunately, the only difference between personal and business financial decisions is the lingo. The financial principles are identical so that everything we've learned thus far can be applied to any business. That said, the lingo can be daunting and the settings foreign at first. So, let's set the stage here by providing some context for how a business is viewed from a financial perspective.

This chapter

- explores what a firm is from a financial perspective,
- introduces some of the connections between personal and corporate finance,
- highlights how the Fundamental Value Relation is behind the decision-making of both the firm and its investors, and
- considers how objectives other than financial ones may be relevant for companies.

We'll see that a firm is really just a means for allocating cash flows - money comes in and out of a firm all the time. To keep the discussion grounded, let's focus on a specific company, Microsoft Corp., which is visualized in Figure 4.1.

We chose Microsoft because most people are familiar with their products. If you're not, Microsoft makes and sells computer software (e.g., Windows, Office), hardware (e.g., Surface laptops and tablets, Xbox consoles), and services (e.g., Azure). Don't misinterpret this choice as having any particular significance. In other words, our discussion below applies to most every company. The scale and numbers may differ, but the ideas are the same.

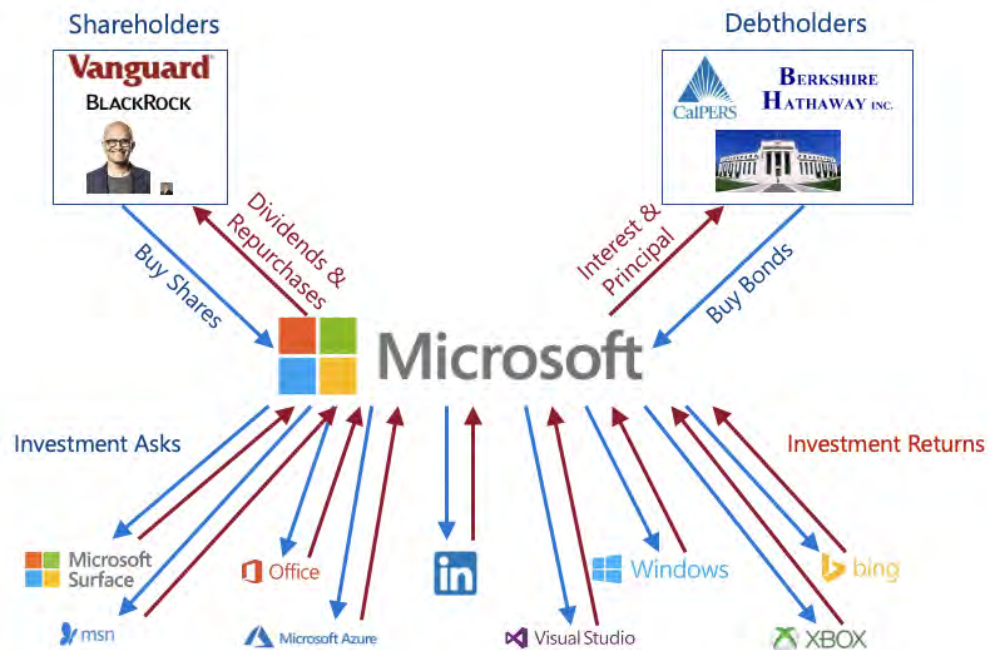


Figure 4.1: Microsoft from a Financial Perspective

4.1 Investors

Investors are the financial starting point for any company. They provide the **capital**, i.e., money, and sometimes resources, such as labor, that companies need to produce, distribute, and sell goods and services. Broadly speaking, there are two types of investors - debtholders and shareholders - both of whom give resources (e.g., money, labor) to the company in exchange for the promise of money in the future.

4.1.1 Shareholders

Shareholders (a.k.a., **stockholders**, **equityholders**, **equity investors**, **owners**) are represented by the box on the top left of the figure. While there are many shareholders in

Microsoft, we've listed a few to discuss. Vanguard and BlackRock correspond to large institutional investors that own Microsoft stock. These institutions own shares of stock on their own account, and also act as custodians for the shares of other Microsoft investors. That is, Vanguard and BlackRock hold the shares owned by other investors who hold accounts with these intermediaries.

There are also two photos in the box. The first, relatively large photo is of Satya Nadella, the CEO of Microsoft as of 2021. The second, relatively tiny photo is of me, the author of this book. Satya and I are both Microsoft shareholders because we own shares of the company. Like the institutions, I purchased Microsoft stock. Satya received shares from the company in exchange for his work. Satya owns many more shares than I do, hence the size difference in our photos. Nonetheless, like Vanguard and BlackRock, Satya and I are shareholders, and therefore owners, of the company.

The blue arrow labeled “Buy Shares” and pointing from the Shareholders box towards Microsoft corresponds to resources leaving shareholders and going to Microsoft. And, with the exception of employees like Satya, this means a cash outflow for shareholders and a cash inflow for Microsoft. Microsoft has tens of thousands of shareholders, all of whom have provided capital or other resources to the company.¹

Of course, shareholders give money to Microsoft, and other companies, because they expect to receive even more money in the future. The red arrow pointing from Microsoft towards the Shareholders box reveals how shareholders get that money, through dividends and share repurchases. Dividends are periodic distributions to shareholders that consist of either cash or more shares of the company. Share repurchases occur when the company buys back shares on the open market, though not everyone has to sell their shares. Both dividends and share repurchases are ways in which companies return money to their shareholders and, as such, are critical to getting money from investors in the first place.

4.1.2 Debtholders

Debtholders (a.k.a., **creditors**, **lenders**) are represented by the box on the top right of the figure. These investors have lent money to Microsoft. In that box, we have the

¹To be precise, I, and perhaps Vanguard and BlackRock, purchased shares on an exchange, like the NASDAQ, from an anonymous seller in what's called the **secondary market**. However, Microsoft had to issue those shares to someone initially on the **primary market** and when it did, it received a bunch of money. The primary market is typically limited to institutional investors or very high net worth individuals (i.e., rich people). So, while I didn't directly give money to Microsoft, I did so indirectly by buying shares from someone else who may have.

logos for CalPERS and Berkshire Hathaway. CalPERS, which stands for the California Public Employees' Retirement System, manages the retirement savings for California public employees. As of 2021, CalPERS was responsible for managing \$469 billion of savings for 1.5 million people. As such, their primary job is wisely investing those savings and one of their investments is in Microsoft bonds, which are just loans to the company. Berkshire Hathaway is the company founded and run by Warren Buffet, the famous U.S. investor. Berkshire owned Microsoft bonds as of 2021. Finally, the building in the photo is that of the Federal Reserve, the central bank of the United States, which also owned Microsoft bonds as of 2021.

The blue arrow labeled “Buy Bonds” and pointing from the debtholders box towards Microsoft corresponds to resources leaving debtholders and going to Microsoft. This means a cash outflow for debtholders and a cash inflow for Microsoft. Like shareholders, Microsoft has many creditors. Some own bonds. Others, such as banks and hedge funds, own loans. Regardless, all debtholders have lent money to Microsoft.² And, like shareholders, debtholders give money to Microsoft and other companies because they expect to receive even more in the future. Debtholders receive this money through interest and principal payments, as indicated by the red arrow pointing from Microsoft towards the debtholders box.

There is no guarantee that investors - shareholders and debtholders - will get more than they gave. In fact, there is no guarantee that they will get any money back. Because there is no guarantee, investing in Microsoft or any other company is risky.

4.1.3 Fundamental Value Relation

The Fundamental Value Relation is central for investor decision-making. Shareholders use this relation to estimate the value of company stock. In doing so, they can compare their valuation to the market price to determine if the shares are over-, under-, or fairly priced. The cash flows, $CashFlow_t$, consist of future dividends and capital gains (or losses) when they sell. The discount rate, r , for these cash flows is the **expected stock return** or **levered return**, which is what stock investors expect to earn on each dollar they invest. This same discount rate is also called the **equity cost of capital** or **levered cost of capital** by the firm. What is a benefit to an investor is a cost to the firm that must pay for the investor's return. In other words, there are two sides to r - the return investors expect to receive and the cost to companies of raising money. Ultimately, these two sides are talking about the same thing.

²Much like how I purchased shares on an exchange, these and other institutions may have purchased bonds or loans in **over the counter (OTC)** markets from other debtholders.

Similarly, debtholders use the Fundamental Value Relation to determine the fair price for their loans. Debt cash flows consist of future interest and principal payments. The debt discount rate is the **expected debt return**, which is what debt investors expect to earn on each dollar they invest. The debt discount rate is also called the **debt cost of capital** by the firm. Again, the investment earnings to creditors are costs of financing paid by the firm.

4.2 Business Decisions

When Microsoft receives money from investors, it allocates or spends it by considering **investment asks** or **requests for funds** from its employees. This process of allocating money is referred to as **capital budgeting**. Exactly how firms should allocate their money is covered in chapters 5 and 6. The key implication of this process for our purposes here is that capital budgeting reflects a cash outflow from Microsoft, as represented by the blue arrows labeled “Investment Asks” and pointing towards the Microsoft products and services (MSN, Office, Azure, Windows, Xbox, etc.)

Microsoft takes the money it receives from investors and spends it in a manner that it believes will generate even more money in the future. The money that’s generated from Microsoft’s spending represent the investment returns indicated by the red arrows pointing away from the products towards Microsoft. What does this mean more simply? Microsoft is spending money all the time on its employees, buildings, server farms, utilities, etc. so that it can sell products and services. In other words, Microsoft is generating investment returns through its sales, cost savings, and other channels we’ll explore later.

As the business receives money from its investments, it has another decision: What to do with the investment returns. It could return this money to its investors - pay dividends to or repurchase shares from its shareholders, or pay interest and principal to its debtholders. Microsoft could also reinvest (i.e., spend) the money in existing or new projects. Finally, the company could save the money for future investment or future distributions to its investors. These decisions are, or at least should be, governed by what is best for the shareholders who own the company.

4.2.1 Fundamental Value Relation

The Fundamental Value Relation is central for business decision-making. Businesses use this relation to assess the financial viability of individual projects, and to choose among competing projects. The cash flows consist of the sales less expenses and investments of

projects. The discount rate is the cost of capital for projects, which is determined by the expected returns its investors demand and how the company finances the project - debt, equity, or a mix. Critically, we'll see that companies, just like us, should use the NPV criterion to make their decisions.

4.3 Nonfinancial Considerations

The traditional view of business, often attributed to the economist Milton Friedman, is that corporations should make decisions that maximize value to their shareholders - i.e., the owners of the company. This begs the question: Is there more to business than making money? There is a great deal of debate over this question both in and out of academia.

Many point to **Corporate Social Responsibility (CSR)** and **Environment, Social, and Governance (ESG)** initiatives as evidence that companies are concerned about more than just profits. These initiatives ask business to make decisions with other, possibly nonfinancial concerns in mind. What are these concerns? Employee welfare. Environmental treatment. Equal opportunity for everyone. And many others. For example, in 2018 Google reduced the energy consumption of their data centers by 50% when compared to competitors' energy usage. Netflix offers employees a full year of paid parental leave. Starbucks pledged to hire 25,000 US military veterans and spouses. Wells Fargo donated \$6.25 million to support the Covid-19 pandemic response. There are many other examples of initiatives whose financial motives are at best unclear, and most likely absent.

A skeptic might view these initiatives as guises to cut costs or increase revenue. That is, they are really just another form of investment aimed at maximizing shareholder value. For example, hiring veterans or donating money to support pandemic relief can be viewed as a form of advertising and an attempt to garner goodwill among potential customers. Likewise, Google's energy usage reduction almost surely came with substantial cost savings in reduced utility bills.

The challenging question is what to do if an investment increases the welfare of one set of stakeholders (e.g., employees, city residents) at the expense of investors? Some investors may like this idea because they value doing good more than they value money. Others may not, perhaps preferring to do good on their own. In the latter case, objectives other than investor wealth maximization may not be viable long-term strategies because most investors have historically given companies money in order to get more money in return. What's the right balance? Companies, investors, employees, customers, governments, etc. are continuously searching for that balance.

4.4 Key Ideas

A firm is kind of like a risky, money processing machine. (We don't want to say money laundering machine.³) Investors give the firm money in the hope it will return even more money to them in the future. The firm takes investors' money and spends it on the goods and services it believes will generate even more money. The firm then takes the earnings from its investments and returns some to its investors and reinvests the rest to generate more money. Thus, the firm in some sense is really just a means for creating more money. That said, this discussion raises a lot of questions.

- How exactly should companies allocate their resources?
- What does it mean to be a shareholder? Debtholder? To what are they entitled, and what rights do they have?
- How can we implement our Fundamental Value Relation to estimate stock prices? Bond prices?
- From where do the expected returns on stocks and bonds come? I.e., How much should investors expect to earn when they buy a stock or bond?
- What happens to investors' returns if they invest in lots of different stocks and bonds?

The remainder of this section is aimed at answering these questions. Before diving in, hopefully we can start to see some of the connections between what's been discussed in previous chapters focusing on personal finance and what's to come in subsequent chapters focusing on corporate finance.

Investors in companies - shareholders and debtholders - are just savers! Remember our retirement savings program in which we earned a return on our savings while working and in retirement? That return often comes from the companies in which we invest. We give these companies money sometimes in the form of a loan (e.g., bond), sometimes in exchange for stock. Companies take our money and invest it in projects. If those projects pay off, the companies return to us even more money than what we gave them.

Likewise, we make capital budgeting decisions all the time, such as whether to buy a house, car, computer, etc. We also have to decide how to finance these purchases. Do we

³**Money laundering** refers to the process of converting illicit funds into seemingly legal funds by hiding the origin of the funds through a series of transactions that mix “dirty” money with “clean” money.

pay with cash? Credit (e.g., mortgage, auto loan, credit card)? A combination? Companies are doing the same thing. They are making capital budgeting decisions, such as whether to buy a plant, piece of property, equipment, etc. And, they too have to decide how to finance these purchases: Cash? Debt? Stock?

What is critical to keep in mind going forward is that despite the different lingo, jargon, and acronyms, the **financial principles we use to make personal finance decisions are the same as those we use to make corporate finance decisions**. There is only one “finance” that is based on our Fundamental Valuation Relation and it doesn’t care about the context, only the cash flows (*CashFlow*) and discount rate (r).

Chapter 5

Project Viability

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

This chapter is a first look at business decision making or **capital budgeting**. Specifically, we

- introduce the decision criteria businesses use to determine whether or not to undertake a project,
- show how the cash flows in our Fundamental Value Relation are computed for all business decisions,
- draw explicit links between common key performance indicators (KPIs) on which managers focus (revenue growth, margins, earnings, days receivable/payable, inventory turnover, etc.) and value,
- apply our Fundamental Value Relation to investigate a hypothetical scenario in which Dell Inc. must decide whether to produce and sell a tablet, and
- highlight the advantages and limitations of formal business decision making by performing sensitivity analysis on our tablet project.

We'll limit ourselves in this chapter to considering the most basic decision - determining whether a project or investment is viable by determining whether it creates or destroys value. Subsequent chapters explore extensions including project selection and how to handle constraints, such as budgets or headcount limits. This chapter assumes you are comfortable with basic accounting. If not, read or refer back to appendix B.

5.1 Decision Criteria

Managers use several different criteria for making decisions. The three most popular according to survey evidence are **net present value (NPV)**, **internal rate of return (IRR)**, and **payback period**.¹ While NPV should be familiar, IRR and payback period are new. We'll discuss and illustrate each with several examples beginning with a very simple one. Imagine our company has a sales initiative costing \$100 million today and generating \$50 million of cash flow per year for the next three years starting one year from today. The opportunity cost for the project is 10%. The timeline is presented in Figure 5.1.




Figure 5.1: Sales Initiative Timeline

Before determining whether this initiative is worth pursuing, it's important to introduce the concept of **incremental** cash flows. When assessing a project, we need only focus on the cash flows that arise as a consequence of that decision. That is, the only cash flows that matter for assessing a project are those that are affected by the project. For example, if our company's existing cash flows without the project are projected at \$2 billion per year over the next three years regardless of whether the sales initiative is undertaken, then these existing cash flows are irrelevant for the evaluation of the initiative. As such, our timeline of cash flows does not include the existing cash flows to the company.

If instead, the sales initiative would increase existing cash flows by 1% for the next three years, in addition to generating \$50 million cash flow perhaps from new sales or cost savings created by the project, then we need to consider that (\$20 million) boost in the evaluation of the sales initiative.

¹John Graham and Campbell Harvey show that these three criteria are the most popular decision criteria used by Fortune 500 CFOs in their 2001 study, "The theory and practice of corporate finance: Theory and evidence from the field," *Journal of Financial Economics* 60, 187-243.

The easiest way to determine what's incremental is to simply ask: What is new, different, or changes as a result of the project or our decision? Equivalently, what is the difference between the company before and after the decision? If something changes or is different as a result of our project, then we have to account for it in our analysis. If it doesn't change, then it's irrelevant.

5.1.1 Net Present Value (NPV)

Recall from chapter 1, NPV measures the difference in the present values of the costs and benefits of a decision. When NPV is positive, value is created, when negative, value is destroyed. Mathematically, NPV is defined in equation 5.1.

$$NPV = CashFlow_0 + \frac{CashFlow_1}{(1+r)} + \frac{CashFlow_2}{(1+r)^2} + \frac{CashFlow_3}{(1+r)^3} + \dots \quad (5.1)$$

Often the initial cash flow, $CashFlow_0$, is negative and corresponds to the cost or price of the investment. However, there is no restriction on the signs of the cash flows.

The NPV (\$mil) of our sales initiative is

$$NPV = -100 + \frac{50}{(1+0.10)} + \frac{50}{(1+0.10)^2} + \frac{50}{(1+0.10)^3} = \$24.34 \text{ million.}$$

This project creates \$24.34 million of value for the owners of the company, i.e., shareholders, in today's dollars. That is, the total value of the company's equity increases by the NPV of the project. Thus, this is a viable project according to the NPV criterion.

Another way to look at this calculation is to recall our fundamental value equation 1.2. The *value* of the sales initiative is the sum of the discounted *future* cash flows, or

$$PV = Value_0 = \frac{50}{(1+0.10)} + \frac{50}{(1+0.10)^2} + \frac{50}{(1+0.10)^3} = \$124.34 \text{ million.}$$

In other words, a fair price for this investment is \$124.34 million. However, the actual price is only \$100 million. So, "buying" this investment for \$100 million is a good deal because it creates \$124.34 million in today's dollars.

From where does this value come? Competitive and comparative advantage, technological innovation, barriers to entry, favorable regulation, greater bargaining power with suppliers or customers, etc. are all potential sources of value in projects. The point of a for-profit business is to create value, and NPV measures that value in dollar terms. As such, *the NPV rule should always be used for decision making.*

Before turning to other decision criteria, it's important to emphasize that the NPV criterion used in business and that used in personal finance are the same. In our personal lives, we strive to make positive NPV decisions. Businesses do the same.

5.1.2 Internal Rate of Return (IRR)

While the NPV criterion should always be used in decision making, others are often considered. The **internal rate of return** or **IRR** of a project is the *one* discount rate such that the net present value - equation 5.1 - of the project equals zero. Mathematically, the internal rate of return is the *IRR* such that the following equality holds.

$$NPV = 0 = CashFlow_0 + \frac{CashFlow_1}{(1 + IRR)} + \frac{CashFlow_2}{(1 + IRR)^2} + \frac{CashFlow_3}{(1 + IRR)^3} + \dots \quad (5.2)$$

The IRR is the **break even discount rate**; the discount rate that makes us no better or worse off from undertaking the investment (i.e., $NPV = 0$). The IRR is occasionally referred to as the **return on investment (ROI)**, though ROI often refers to an accounting measure of return.²

When the initial cash flow, $CashFlow_0$ in equation 5.1, is the price of the investment, then the IRR is also the one discount rate such that the value of the investment equals its price (i.e., $NPV = 0$).

$$Price_0 = \frac{CashFlow_1}{(1 + IRR)} + \frac{CashFlow_2}{(1 + IRR)^2} + \frac{CashFlow_3}{(1 + IRR)^3} + \dots$$

Viewed this way, the IRR is the periodic return investors earn from the project assuming that all intermediate cash flows are reinvested at the same IRR.

The **IRR decision criteria** says accept a project when the IRR is greater than the project's cost of capital, reject when it is below. As discussed earlier, the cost of capital for a company is just the return that investors in the company expect to earn. What is a cost to one party - the company - is a benefit to the other - investors. The IRR rule ensures that when companies take investors' money, companies invest that money in projects offering a return at least as high as what their investors are expecting to earn. Intuitively, this makes sense. The company takes investors money, which costs r , and invest it at a higher rate, *IRR*, thereby creating value.

Consider our sales initiative. Using a computer, we can solve the following equation to get the IRR.³

$$NPV = 0 = -100 + \frac{50}{(1 + IRR)} + \frac{50}{(1 + IRR)^2} + \frac{50}{(1 + IRR)^3} \implies r = 0.2338$$

²ROI can refer to return on invested capital (ROIC) and return on equity (ROE). The former is the ratio of net operating profit after taxes to total book capitalization (debt plus equity). The latter is the ratio of earnings to book equity.

³In Excel or Google Sheets, the function *IRR* will return the IRR for a set of cash flows. See Appendix C for details on this function and its use.

The IRR of our sales initiative is 23.38% per year. In other words, investing \$100 million in this project today will generate a 23.38% return on our investment each year assuming we can reinvest each intermediate cash flow - years 1 and 2 - at the same 23.38%.

If our project's cost of capital is 10%, meaning investors expect 10% per year in return for their money, then the project returns more than double what our investors expect. So, we would accept this project on the basis of the IRR criterion. If the IRR was less than our project's cost of capital, we would reject the project.

The threshold return that a project's IRR must clear is often referred to as a **hurdle rate**. This hurdle rate is sometimes equal to the cost of capital, sometimes greater than it. The rationale for choosing a hurdle rate greater than the cost of capital is not well-founded. One plausible reason is that because companies have to estimate their cost of capital a higher hurdle rate mitigates the risk of accepting bad projects. For example, if the true cost of capital is 10%, but the company estimates it to be 6%, then projects with IRR's between 6% and 10% will be accepted when they should be rejected.

A more problematic rationale for a hurdle rate greater than the cost of capital is the assumption that projects with higher IRRs are more valuable. This need not be true. What will be true is that projects with higher IRRs tend to be riskier on average. So, an unintended consequence of inflating our high hurdle rate is an increase in the risk of the business as we focus on riskier projects.

As we'll see in chapter 6, the IRR suffers from several shortcomings that limit its use especially when comparing projects. However, one requirement that the IRR rule must satisfy even when evaluating projects by themselves is adherence to the **sign rule**, which requires all cash outflows (negative cash flows) precede all cash inflows (positive cash flows). When this condition is true, the IRR rule leads to the same decision as the NPV rule ensuring we make the correct decision as it relates to financial value. Table 1 gives examples of cash flow sign patterns in which the NPV and IRR decision rules will lead to the same decision (left column) or possibly different decisions (right column).

We need to check the signs of our project's cash flows before using the IRR rule. In our sales initiative example, the signs of these cash flows are -,+,+,+ corresponding to the \$100 million outflow today followed by the \$50 million inflow each year thereafter. Because all of the outflows (negative cash flows) precede all of the inflows (positive cash flows), using the IRR rule to assess the sales initiative will lead us to the same conclusion as using the NPV rule.

Instead, imagine the cash flows of our sales initiative were -\$50, \$65, -\$20, and \$100

NPV and IRR lead to...	
same decision	possibly different decisions
-,+,+,+,+	+,-,+,-,+
-,-,+,+,+	+,+,+,+,-,-
-,-,-,-,+	+,+,+,+,-

Table 1: Examples of Cash Flow Sign Patterns in which NPV and IRR Decision Rules Produce the Same or Possibly Different Decisions

million ,perhaps because we had additional investment to make two years from today. The signs of these cash flows are $-,+,-,+$, which do not adhere to the sign rule. There is an outflow, \$20 million *after* an inflow, \$65 million. Using the IRR rule to assess the project with these cash flows may or may not lead to a decision that coincides with the decision implied by the NPV rule. We'll see more examples later.

5.1.3 Payback Period

The **payback period** is the time required to recover the cost of the project. The **payback period decision criteria** selects projects whose payback period falls below a predetermined cutoff (e.g., 2 years). To compute the payback period, we cumulate the project cash flows and identify the time at which these cumulative cash flows switch from negative to positive. For our sales initiative, the computations are in Table 2.

	Period			
	0	1	2	3
Cash flows	-100	50	50	50
Payback Period Computation				
Cumulative cash flows	-100	-50	0	50
Discounted Payback Period Computation				
Discounted cash flows	-100	45.45	41.32	37.57
Discounted cumulative cash flows	-100.00	-54.55	-13.22	24.34

Table 2: Sales Initiative (Discounted) Payback Period Calculations

The cumulative cash flow in period 0 is equal to that period's cash flow because there are no prior cash flows. The cumulative cash flow in period 1 is the previous period's cumulative cash flow, -100, plus the current period's cash flow, 50, which equals -50. The same logic applies for future periods; hence period 2's cumulative cash flow is $-50 + 50 = 0$.

The payback period is equal to the period in which the cumulative cash flows equal zero - two years in our sales initiative example. When the cumulative cash flows do not equal zero at the end of a period, practitioners will approximate the payback period by “eyeballing it,” or using linear interpolation. See the Technical Appendix to this chapter for details of the latter.

The decision to accept or reject this project depends on our cutoff. If the cutoff is one year, we would reject our sales initiative because the payback period is after the cutoff. The project takes too long to payback. If the cutoff is three years, we would accept our sales initiative because the payback is before the cutoff.

One concern with the payback period is that it doesn’t account for the time value of money. The **discounted payback period** addresses this shortcoming by using discounted cash flows. The bottom two rows of the table 2 present the discounted and cumulative discounted cash flows. The discounted payback period is a little over two years, longer than the payback period.

The payback period and its cousin the discounted payback period are appealing decision criteria for several reasons. First, it is easy to compute and easy to communicate. Second, it conveys important information about the project’s affect on corporate liquidity or cash resources. Longer payback periods expose companies to greater risk, and require financial managers to find money from other sources while the project evolves.

The payback period also suffers from several shortcomings, the most obvious of which is that it has nothing to say about the value of a project. How long it takes to recover an investment - measures in time - can be completely disconnected from value. Related, the threshold for what distinguishes acceptable from unacceptable projects is arbitrary and removed from value considerations, despite often being justified as based on experience or product life-cycle considerations. Because the payback period ignores all cash flows after the threshold, the timing of the cash flows dominates the decision criteria. As a consequence, decision makers will select projects that payoff more quickly, as opposed to projects that are more valuable, leading to myopic decisions that may stifle long-term growth and destroy value.

Table 3 illustrates some of these problems with four projects, each having a 10% cost of capital. Project 3 creates the most value, but has the longest payback period and would not be accepted if the threshold were less than 2.2 years. Project 4 has the quickest payback period but destroys value because of costs at the end of the project. Project 2 has a similarly quick payback, but only because of its front-loaded cash flow. Project 1 offers much more value than Project 2 despite its longer payback period.

	Cash Flows				Payback Period (Years)	NPV
Project 1	-100	70	50	40	1.8	35.01
Project 2	-100	115	0	0	0.9	4.55
Project 3	-100	30	50	120	2.2	58.75
Project 4	-100	130	-40	0	0.8	-28.51

Table 3: Shortcomings of the Payback Period Criterion

5.1.4 Money Multiples

While we'll focus on NPV, IRR, and payback criteria in this book, it is worth mentioning a few other criteria used in practice. Unfortunately, their use and definition differ across practitioners so it's always a good idea to clarify how exactly each is being calculated.

Cash-on-cash return or **cash yield** is computed as the ratio of the periodic cash received to the total cash invested. Some practitioners use a pre-tax measure of cash flow, others a post-tax measure. Ultimately, the cash-on-cash return tells us the amount of money generated per dollar invested. Like returns, cash-on-cash returns tend to reflect the risk of the investment. In our sales initiative, the annual cash return is $50/100 = 50\%$. When the cash flows are not constant, the cash return will vary from period to period. While informative, a cash return is incomplete, focusing on just one cash flow at a time and therefore not a terribly useful decision metric on its own.

A **cash multiple** or **multiple on invested cash (MOIC)** is the ratio of the *total* cash received to the total cash invested. In our sales initiative example, the cash multiple is $(50 + 50 + 50)/100 = 1.5x$. This result is read "1.5 times" and means we are getting 1.5 times any money we invest. Unlike the cash-on-cash return, the cash multiple accounts for *all* of the cash flows we will receive from our investment. However, this metric ignores the timing of the cash flows and the time value of money. Table 4 illustrates how different projects with identical cash multiples can have very different NPVs. Thus, this metric is really only useful when comparing projects with similar time horizons, risk profiles, and payouts, and even then is potentially problematic.

Why do practitioners use these metrics despite their known shortcomings? Some may not know of their shortcomings. More likely, inertia and historical practice lead to the continued use of these metrics, which are convenient and a part of the language of finance. Additionally, in many situations, projects may have features and similarities that mitigate some of the shortcomings of these metrics. Nonetheless, while these metrics are useful communication

Projects	Cash Flows (\$mil)				NPV (r=10%)	Cash Multiple
	0	1	2	3		
A	-100	50	50	50	24.34	1.5x
B	-100	150	0	0	36.36	1.5x
C	-100	0	0	150	12.70	1.5x

Table 4: Cash Multiple vs. NPV

device, NPV is unambiguously the best decision criterion of the the ones discussed here.

5.1.5 Summary

Our three primary decision criteria are

1. NPV: Accept positive NPV projects; reject negative NPV projects.
2. IRR: Accept projects whose IRR is greater than the cost of capital; reject projects whose IRR is less than the cost of capital.
3. Payback period: Accept projects whose payback period is less than a threshold time; reject projects whose payback period is greater than the threshold time.

The process we use to get the first two metrics, NPV and IRR, is referred to as **discounted cash flow** or **DCF** analysis. In fact, we've been doing DCF all along - discounting cash flows. However, when people in practice refer to DCF analysis, they are typically referring to applications in corporate settings.

Many academics will say: "Use the NPV criterion and forget the rest." Strictly speaking, they're correct. Practically speaking, other criteria serve a purpose in arguing for or against a project. NPV will lead you to the correct answer, but it can be difficult to communicate, especially to investors who are focused on returns. IRR resonates more with investors, and many managers, and is therefore a useful communication tool. The payback period provides additional information related to risk and liquidity that more finance-oriented executives may find valuable. Thus, a judicious use of all three in business case presentations is probably the best approach. The challenge comes when these criteria don't agree, in which case NPV must be relied upon for decision making. We'll see why later.

An even bigger challenge than dealing with conflicting decision criteria is simply implementing them. Take the sales initiative we've been using. From where did the cash flows

come? The cost of capital? The next section focuses on how to estimate cash flows, taking the cost of capital as given. Why? Because estimating cash flows for projects is relevant for all of us, regardless of our particular function within the company. If we're going to justify spending money on anything, we must present the financial costs and benefits, i.e., the cash flows.

The cost of capital is the purview of finance specialists. Financial executives should tell us what the cost of capital or hurdle rate is *before* we present our business case. In other words, we should know the rules of the game before playing. That said, understanding the cost of capital is important for engaging with finance professionals, as well as for investing, and is discussed in greater detail in later chapters.

5.2 Free Cash Flow

To estimate any of our decision metrics, we need to estimate cash flows. For capital budgeting purposes, cash flows are referred to as **free cash flows** (FCF), and there are two types: (i) **unlevered** and (ii) **levered**. When people say “free cash flow,” they are almost always referring to the former. We'll follow that convention and focus on unlevered free cash flows in this chapter. Levered free cash flows require additional information, are used less frequently in practice, and are covered later in the book.

Free cash flow is the money generated by or needed for a project or an entire company after all revenues, expenses, investments, and taxes have been considered. When free cash flow is positive, companies have extra money that they can do with as they please, such as

- distribute to creditors by paying interest or principal,
- distribute to shareholders by paying a dividend or repurchasing shares,
- invest in other projects,
- distribute to employees in the form of bonuses, or
- retain for future investments.

In other words, positive free cash flow is what firms are trying to generate with their business. The more positive free cash flow a company can generate, the more valuable the company will be, all else equal. That is, the *CashFlow* terms in our our fundamental value equation will be larger.

In order to generate positive free cash flows, companies often first generate negative free cash flows. To make money, we have to spend money. When free cash flow is negative, companies must find money to fund the project or company. This money can come from

- creditors like banks or bondholders,
- shareholders (old or new),
- free cash flows from other projects,
- cash savings, or
- the sale of assets.

Formally, unlevered free cash flow for a period (e.g., year, quarter, month) is defined as

$$\begin{aligned} \text{Unlevered free cash flow} = & (Sales - Expenses - D\&A) \times (1 - \tau) \\ & + D\&A - NLTI - NWCI. \end{aligned} \tag{5.3}$$

Equation 5.3 looks daunting but is simply a representation of how businesses function. The acronyms are as follows: *D&A* is depreciation and amortization, τ is the marginal tax rate, *NLTI* is net long-term investment, and *NWCI* is net working capital investment. Let's discuss each of these components and how they affect cash flow below.

5.2.1 Profits

The first line of equation 5.3,

$$(Sales - Expenses - D\&A) \times (1 - \tau) \tag{5.4}$$

is after-tax operating profits, more frequently referred to as **net operating profit after taxes** or **NOPAT**. NOPAT has a lot of synonyms including: **unlevered earnings (UE)**, **earnings before interest after taxes (EBIAT)**, **unlevered net income (UNI)**, and **net operating profit less adjusted taxes (NOPLAT)**.

NOPAT starts with *Sales*, which is how most projects and business generate cash flows. However, not all successful projects require sales. Cost cutting measures and asset sales can be valuable projects in which no revenue is generated. Rather, reductions in expenses and negative investment drive free cash flows and therefore value in these situations.

From *Sales* we subtract cash expenses, *Expenses*, equal to the sum of cost of goods sold (COGS) and selling, general, and administrative (SG&A) expenses. Cash expenses include salaries, rent, utilities, marketing, production costs, distribution costs, research and development costs, etc. *Sales* minus *Expenses* is **EBITDA** - Earnings Before Interest, Taxes, Depreciation, and Amortization.

After deducting cash expenses, we subtract non-cash expenses, which typically mean depreciation and amortization (*D&A*). In resource extraction industries such as oil and gas and mining, depletion is another common non-cash expense that is deducted from EBITDA.⁴ *Sales* minus *Expenses* and *D&A* is **EBIT** - Earnings Before Interest and Taxes.

Finally, *EBIT* times $(1 - \tau)$ equals NOPAT, where τ is the appropriate tax rate. When evaluating a project, we use the **marginal tax rate** that measures the tax paid on each incremental dollar. When valuing an entire company, we use the **effective tax rate** or **average tax rate** applying to all profits of the company. Let's annotate our free cash flow expression (equation 5.3) with the different earnings measures.

$$\text{Unlevered free cash flow} = \underbrace{\left(\underbrace{\overbrace{\text{EBITDA}}^{\text{EBIT}} - D\&A}_{\text{NOPAT}} \right) \times (1 - \tau)}_{\text{NOPAT}} + D\&A - NLTI - NWCI.$$

Free cash flow starts with NOPAT, which is built up from measures of operating earnings, EBITDA and EBIT.

Notice in our expression for NOPAT that we didn't subtract interest expense or add interest income. More generally, we didn't adjust NOPAT for any non-operating expenses or income. This exclusion means NOPAT will typically be different from net income on the bottom of an income statement. Table ?? presents the difference between NOPAT and net income using Microsoft Inc.'s 2021 income statement. The other expenses (income) account is ignored in the computation of NOPAT, and the taxes are computed by applying the tax rate 13.8% to EBIT instead of Pre-tax income.

This difference between NOPAT and net income means that how much interest the company pays on what it borrows or earns on what it saves does *not* affect unlevered free cash flow. Put differently, the company's **financial policy** - how it chooses to fund projects

⁴Strictly speaking, we should subtract *all* non-cash expenses and add all non-cash revenues that affect the tax calculation. Then we should add back the non-cash expenses and deduct the non-cash revenues.

	2021		2021
Sales	168,088	Sales	168,088
Cost of sales	40,546	Cost of sales	40,546
Gross profit	127,542	Gross profit	127,542
SG&A	45,940	SG&A	45,940
EBITDA	81,602	EBITDA	81,602
Depreciation & Amortiation	11,686	Depreciation & Amortiation	11,686
EBIT	69,916	EBIT	69,916
Other expenses (income)	(1,186)	Taxes	9,667
Pre-tax income	71,102	NOPAT	60,249
Taxes	9,831		
Net income	61,271		

Table 5: Microsoft Inc. Income Statement, Net Income, and NOPAT

or the business as a whole - does not affect unlevered free cash flow. In this sense, unlevered means unaffected by leverage or debt. Where our CFO gets money to pay for our project - loans, cash savings, stock issuances - is has no impact on our calculation of unlevered free cash flow. This feature of free cash flow allows us to assess the financial costs and benefits of a project, i.e., free cash flows, without having to worry about how it will be paid for, which is the CFO's job.

This freedom does *not* mean financing is irrelevant. On the contrary, how the project is financed can create significant value, which is in large part why CFOs are paid so much! A big part of their job is to find the cheapest financing possible, which means raising money with the lowest cost of capital, r . Everyone else's job is to find projects with really big free cash flows. This is how value is maximized by firms. Everyone is searching for ideas that will generate large cash flows while the CFO focuses on finding cheap money with a low cost of capital. Large cash flows and a low cost of capital mean high value according to the Fundamental Value Relation (equation 1.2).

5.2.2 Investments

In addition to ignoring interest expense, free cash flow differs from earnings because of the second line in equation 5.3,

$$D\&A - NLTI - NWCI$$

To arrive at free cash flow, we first add back depreciation and amortization expense, $D\&A$, which is a non-cash expense. When an asset is depreciated or amortized by accountants, no money leaves the company. This fact begs the question: If depreciation and amortization

doesn't correspond to money leaving the company, why is it in our free cash flow expression? To understand why, let's distribute the $(1 - \tau)$ term in equation 5.3 and focus on where $D\&A$ appears.

$$\begin{aligned} -D\&A \times (1 - \tau) + D\&A &= -\cancel{D\&A} + D\&A \times \tau + \cancel{D\&A} \\ &= D\&A \times \tau \end{aligned}$$

Even though $D\&A$ appears twice (three times actually) in equation 5.3, after canceling terms we're left with just $D\&A \times \tau$. This term tells us for each dollar of depreciation and amortization expense, free cash flow *increases* by τ , the project tax rate. Depreciation and amortization expense creates a **tax shield** for the company. By depreciating and amortizing assets, the company reduces its taxable income and, as such, reduces its taxes. From a cash flow perspective, paying less is equivalent to receiving more.

Net Long Term Investment (NLTI)

The next term, **net long-term investments** ($NLTI$), captures money spent on or received from long-term assets on the balance sheet. For example, **capital expenditures** correspond to money spent on plant, property, and equipment, PP&E for short. NLTI also includes money spent on intangible assets, such as patents, trademarks, licensing agreements, and software. Finally, NLTI could include acquisitions of whole companies. The “net” in this term refers to the possibility that some projects may require the sale of assets, in which case $NLTI$ could be negative. Because $NLTI$ is preceded by a minus sign in our free cash flow formula, any asset sales - negative $NLTI$ - result in increases to free cash flow.

Many companies debate whether to “OpEx or “CapEx” money they spend. In other words, companies will sometimes have the choice, under accounting rules, to classify money they spend as an operating expense that gets reported on the income statement or as a capital expenditure (or intangible investment) that gets reported on the balance sheet. The tradeoff in this decision is between maximizing the value of a tax shield and minimizing the impact on earnings.

For example, imagine we are facing a 21% tax rate, and we spend \$10 million on design and engineering plans for a new production facility. If we OpEx this spend, then we will reduce our income tax by $10 \times 0.21 = \$2.1$ million this year. But, our earnings will be reduced by \$7.9 million. If we CapEx this spend, our earnings are unaffected this year, and we'll pay \$210,000 more in taxes than we would had we treated it as OpEx. These capital expenditures will gradually find their way onto the income statement starting next year, at which point we

will receive a (smaller) tax shield each year over the life of the investment, but our earnings will be much smoother as we spread the cost over several years in the future.

It may seem that classifying any spend as operating expense, in so far as it is allowable under accounting rules, would be the obvious choice because it unambiguously creates value through a larger tax shield. However, because corporate managers typically have more information than investors, earnings are often an important signal of the companies health and future prospects. A sharp decline in earnings may lead to a reduction in stock prices and corresponding increase in a company's cost of capital, r , if the executive team cannot clearly communicate to investors the reason for the decline. Additionally, if the company has no taxable income to shield, perhaps because it is losing money, then the operating expense today would have no impact on taxes today. Thus, the choice between OpEx and CapEx is one with potential value implications.

Net Working Capital Investment (NWCI)

The final term in our free cash flow expression is **net working capital investment** (*NWCI*), also known as the **change in net working capital** or sometimes just the **change in working capital**. Net working capital is the difference between current assets (e.g., inventory, accounts receivable) and current liabilities (e.g., accounts payable), both of which are found on the balance sheet. Net working capital are the resources companies rely upon for the day-to-day and short-term operation of the business.

Net working capital *investment* is the periodic change in net working capital. For example, if net working capital was \$300 at the start of the year and \$500 at the end of the year, the investment in net working capital is $500-300=\$200$ - a cash outflow. This investment could correspond to a build up of inventory, extension of credit to its customers (accounts receivable), or credit from its suppliers (accounts payable). It's useful to think of net working capital investment as the short-term counterpart to net long-term investment. Companies not only have to invest for the long-term, but also for the short-term to keep the business running day-to-day.

Two subtle points regarding net working capital investment. First, we exclude financing in our calculation of net working capital investment. This means short-term debt and long-term debt due within a year, both of which are current liabilities, are not considered a part of net working capital for the purpose of computing free cash flow. Remember, unlevered free cash flows are unaffected by financial policy. Second, projects sometimes require cash be set aside for working capital purposes (e.g., pay salaries and bills). The cash that is required

for these purposes should be included as a current asset in the calculation of working capital. All **excess cash**, or cash not needed for running the company, should be excluded from net working capital. Excess cash could correspond to profits awaiting distribution to investors.

5.2.3 Summary

Equation 5.3 provides a recipe for computing the free cash flow for *any* project over any period (e.g., quarter, year). Different projects in different industries all rely on this same equation to compute free cash flow. This is important. The framework is the always the same, only the numbers differ.

Projects that require no investment in long-term assets or working capital have zeros in the place of *NLTI* and *NWCI*. Cost saving initiatives with no sales implications have zeros for *Sales* and negative amounts for *Expenses*. Projects at companies located in countries with no corporate taxes (e.g., Anguilla, Cayman Islands, Jersey) have a tax rate, τ , of zero.

Equation 5.3 also provides a useful framework for articulating the financial rationale for an idea. Each cost or benefit of a project must impact one of the components of free cash flow to affect a decision's value. In other words, when motivating a business decision, it is useful to discuss strategic and operating considerations in light of their impact on revenue, expenses, taxes, or investment, as this is how financial value is created or destroyed.

5.3 Assessing a Stand Alone Project

Let's look at a more involved albeit hypothetical example. The date is January, 2012, and Dell Inc. is deciding whether to develop, produce, and distribute a tablet to compete with the Apple iPad. The project management team has decided to look out over a three year horizon to understand the short-term value of the opportunity and its impact on the firm's financials.

Why three years? Because Dell anticipates a three year product-life cycle. Perhaps they will come to market with another version in three years. Regardless, there is nothing special about three or any other number of years. Our forecast horizon should be long enough to capture all of the costs and benefits of a decision. For example, were we considering the construction of a power plant, our horizon might be 30 or 40 years.

5.3.1 Forecast Drivers

After conferring with experts in the company, the team has identified **forecast drivers**, or assumptions about how the project will affect the finances of the company. These forecast drivers are detailed below and organized in a manner that makes them easy to link back to the free cash flow components.

1. *Sales*

- The market for tablets is forecasted to be 1 million units in 2013, when Dell plans on going to market. The market is anticipated to grow tenfold in 2014 and reach 30 million units in 2015.
- Dell projects a 20% share of the tablet market in 2013, and expects that number to grow by 5 percentage points per year over the following two years.
- Dell's pricing strategy is to enter the market at \$400 per unit and increase the unit price by 10% per year to account for product improvements (e.g., more memory, sharper resolution screen, more apps).
- Dell's laptop division estimates a \$5 million reduction in their revenue when the tablet is launched in 2013 as price cuts will be needed to entice customers to continue buying laptops. This revenue loss is expected to double in each of the following years.

2. *Expenses*

- In 2011, Dell spent \$2 million on focus groups exploring interest in a tablet and what features consumers would like.
- The unit cost of a tablet is \$75 and is forecasted to increase with inflation at 3% per year.
- To support the product, **overhead** expenses including marketing, sales, and administration will cost 60% of annual sales.
- Research and development (R&D) costs are estimated to be \$50 million in 2012, and \$10 million per year thereafter.

3. *NLTI*, Net long-term investment and depreciation

- Dell will need to spend \$400 million in 2012 to retrofit an existing production facility for tablet purposes. The facility will have a usable life of 10 years at which time its salvage value will be \$10 million.

- Three years from today, Dell anticipates selling the facility for \$200 million, what we refer to as a **liquidation value**, to move any future production offshore. *Salvage value is the value at the end of the usable life of the asset and is used to estimate the periodic depreciation (or amortization). Liquidation value is the value when the asset is sold, i.e., how much pre-tax money we'll receive when we sell the asset.*

4. *NWCI*, Net working capital investment

- The treasury department estimates 14 days to collect payments from customers, most of whom will pay by credit card, and 45 days to pay suppliers.
- Dell will pre-build 25% of the following year's unit sales to avoid stock outs, i.e., running out of tablets.

5. τ , Tax rate

- Dell's marginal tax rate for the tablet project is 21%

6. r , Cost of capital

- Dell's cost of capital for the tablet project is 12%.

Figure 5.2 presents the forecast driver section of the spreadsheet used to perform the capital budgeting exercise. This section contains all of the assumptions needed to estimate free cash flows and, ultimately, the decision metrics: NPV, IRR, and payback period. We've organized the forecast drivers to mimic the free cash flow relation for ease of reference.

The different font colors correspond to different cell contents - blue indicates a hard-coded entry, black indicates a formula. The top row identifies the end of the period - year in this example - of the project's life. As always, period 0 is today, January of 2012. The "Step" column controls the annual change our forecast drivers. For example, Dell's market share increases by 5% each year, while its unit prices grows by 10% per year.

5.3.2 Incremental Earnings

Figure 5.3 presents the **pro forma**, or projected, P&L (a.k.a., profit and loss statement) for the tablet project. It shows the incremental sales, expenses, taxes, and earnings of the project at each point in the project's life.

Forecast drivers	0	1	2	3	Step
Sales					
Market size (million units)		1.0	10.0	30.0	
Dell market share		20.0%	25.0%	30.0%	5.0%
Unit price (\$)		400.0	440.0	484.0	10.0%
Loss of laptop revenue (\$mil)		5.0	10.0	20.0	100.0%
Expenses					
Unit cost (\$)		75.0	77.3	79.6	3.0%
SG&A (% of sales)		60.0%	60.0%	60.0%	
R&D (\$mil)	50.0	10.0	10.0	10.0	
Net long-term investment					
Capital expenditures (\$mil)	400.0				
Salvage value (\$mil)				10.0	
Liquidation value (\$mil)				200.0	
Usable life (years)	10.0				
Net working capital investment					
Days receivable	14.0	14.0	14.0	14.0	
Inventory (% next year output)	25.0%	25.0%	25.0%	25.0%	
Days payable	45.0	45.0	45.0	45.0	
Tax rate	21.0%	21.0%	21.0%	21.0%	
Cost of capital	12.0%	12.0%	12.0%	12.0%	

Figure 5.2: Dell Tablet - Forecast Drivers (“Step” dictates how the drivers change over time)

Starting at the top, tablet sales are computed as the product of (i) market size, (ii) Dell’s market share, and (iii) the unit price. Multiplying market size by market share tells us how many units Dell is expecting to sell. Multiplying the number of units Dell will sell by the price per unit tells us the dollar revenue from selling tablets. For example, in its first year after going to market, Dell expects to sell 20% of all tablets sold that year or 0.20×1.0 million = 200,000 units. Each unit will sell for \$400 for total revenue of $200,000 \times 400 = \$80$ million.

The dollar loss in laptop sales (cannibalization) are deducted from the tablet sales to get the net sales for the project. This deduction is important to avoid overstating the benefits of the project. If the tablet isn’t sold, then Dell wouldn’t experience the decline in laptop sales. It is only because of the tablet that revenue from the laptops will decline and, therefore, we have to account for this decline in our analysis. Recall the discussion earlier. To determine what is incremental to a project, we must determine what changes as a result of the project. Anything altered by the project must be accounted for in our cash flow calculation. Put

P&L (\$mil)	<u>0</u>	<u>1</u>	<u>2</u>	<u>3</u>
Sales, tablet	\$0.0	\$80.0	\$1,100.0	\$4,356.0
Cannibalization, laptop sales	0.0	(5.0)	(10.0)	(20.0)
Net sales	0.0	75.0	1,090.0	4,336.0
COGS	0.0	(15.0)	(193.1)	(716.1)
Gross profit	0.0	60.0	896.9	3,619.9
SG&A	0.0	(48.0)	(660.0)	(2,613.6)
R&D	(50.0)	(10.0)	(10.0)	(10.0)
EBITDA	(50.0)	2.0	226.9	996.3
Depreciation and amortization		(39.0)	(39.0)	(39.0)
EBIT	(50.0)	(37.0)	187.9	957.3
Taxes	10.5	7.8	(39.5)	(201.0)
NOPAT	(39.5)	(29.2)	148.4	756.3

Figure 5.3: Dell Tablet - P&L

differently, **capital budgeting exercises must capture *all* of the incremental effects associated with the project - direct and indirect.**

In contrast, **allocated costs** are typically not included in a DCF analysis. Accountants often allocate **overhead costs** - rent, utilities, executive salaries, etc. - to different projects. Just because an accountant allocates a cost to our project does *not* mean that it should be accounted for in our DCF analysis. If those costs exist regardless of whether or not we undertake our project, then they are irrelevant for our analysis. However, if we have to hire a new executive or our utility bill increases because of our project, then we must account for these changes.

Returning to our tablet example, cost of goods sold (COGS) consists of the direct costs in producing the good or service, in this case the cost per unit times the number of units sold. So first year sales of 200,000 units will cost Dell $200,000 \times 75 = \$15$ million.

Selling, general, and administrative expenses (SG&A) consists of the indirect costs in producing the good or service, often referred to as overhead. SG&A is 60% of the forecasted tablet sales, as opposed to net sales. This assumes that there is no reduction in overhead costs associated with the decline in laptop sales, which seems reasonable since the reduction is coming from a price cut as opposed to reduced volume. Annual dollar R&D expenses are given.

Depreciation and amortization are a consequence of the capital expenditures on the production facility. The \$400 million investment is paid for today and begins depreciating the following year. The annual depreciation expense is

$$\text{Annual depreciation} = \frac{\text{Cost} - \text{Salvage value}}{\text{Usable life}} = \frac{400 - 10}{10} = \$39 \text{ million.}$$

Taxes are 21% of EBIT and result in what look like a tax refund - positive taxes - in the first two years of the project. This does not mean that the tax authority, the IRS in the U.S., will be writing Dell checks. Rather, these figures represent tax shields created by the project's losses in the first two years. The tax shield arises because Dell is a profitable company. The losses from the tablet project will be used to reduce the taxable income for the company as a whole, resulting in lower taxes. Remember, paying less and receiving more are two sides of the same coin.

One final point is that nowhere on the P&L is the \$2 million spent on focus groups. The reason this expense has been ignored is because it occurred *in the past*. As such, it is a **sunk cost** and irrelevant for our decision. Including sunk costs in financial decisions is so common that doing so is referred to as the **sunk-cost fallacy**, the belief that money already spent is relevant for a decision. The information gleaned from the focus groups is surely relevant for our decision as it provides information about future sales and other value drivers. However, the \$2 million spent obtaining the information is irrelevant because nothing can be done to change that spend. It is spilt milk, water under the bridge. The lesson: Don't fall prey to the sunk cost fallacy.

Carryforwards

The losses in the first two years of the project generated tax shields because Dell was profitable at the time and had taxable income it could shield. But, what if Dell wasn't profitable? What if Dell had no taxable income? With no taxable income, Dell wouldn't pay any taxes so the project losses wouldn't reduce any taxes, at least not immediately. In this case, Dell could use **operating loss carryforwards** to shield *future* profits from taxes. These carryforwards and corresponding tax implications is presented in Figure 5.4.

In period 0, the \$50 million operating loss (-\$50 million EBIT) is added to Dell's bank of carryforwards. The \$37 million loss in period 1 increases Dell's total carryforwards to \$87 million. The \$87 million of loss carryforwards are used up by the \$187.9 million EBIT in period 2, but taxable income is reduced from \$187.9 million to \$100.9 million. As a result,

Net operating loss (NOL)	0	1	2	3
Start NOL carryforward	0.0	50.0	87.0	0.0
Change in NOL	50.0	37.0	(87.0)	0.0
End NOL carryforward	50.0	87.0	0.0	0.0
Taxable income	0.0	0.0	100.9	957.3
Taxes	0.0	0.0	21.2	201.0

Figure 5.4: Dell Tablet - Operating Loss Carryforwards

the tax bill in year 2 is lower than without the carryforward - \$21.2 million vs. \$39.5 million. The impact of operating loss carryforwards on the P&L can be seen in Figure 5.5.

P&L (\$mil, carryforwards)	0	1	2	3
Sales, tablet	\$0.0	\$80.0	\$1,100.0	\$4,356.0
Cannibalization, laptop sales	0.0	(5.0)	(10.0)	(20.0)
Net sales	0.0	75.0	1,090.0	4,336.0
COGS	0.0	(15.0)	(193.1)	(716.1)
Gross profit	0.0	60.0	896.9	3,619.9
SG&A	0.0	(48.0)	(660.0)	(2,613.6)
R&D	(50.0)	(10.0)	(10.0)	(10.0)
EBITDA	(50.0)	2.0	226.9	996.3
Depreciation and amortization	0.0	(39.0)	(39.0)	(39.0)
EBIT	(50.0)	(37.0)	187.9	957.3
Taxes	0.0	0.0	(21.2)	(201.0)
NOPAT	(50.0)	(37.0)	166.7	756.3

Figure 5.5: Dell Tablet - P&L with Operating Loss Carryforwards

With no taxable income to shield, there are no taxes paid (or reduced) in periods 0 and 1. In period 2, taxes are reduced because of the use of the carryforwards. In the last period of the project (period 3), there are no more carryforwards and the entirety of the project's operating income - \$957.3 million - is taxed at the marginal tax rate.

5.3.3 Net Long-Term Investments

Figure 5.6 presents a depreciation schedule detailing the net long-term investment and corresponding depreciation. Dell spends \$400 million in capital expenditures today to retrofit

the production facility. Over the next three years, this investment in PP&E will depreciate at a rate of \$39 million per year. (See the calculation above.) At the end of the project, three years later, Dell will receive a \$200 million cash inflow from the sale of the facility - its liquidation value.

Deprecation schedule (\$mil)	0	1	2	3
CapEx	400.0			
Start net PP&E		400.0	361.0	322.0
Depreciation		39.0	39.0	39.0
End net PP&E, book value	400.0	361.0	322.0	283.0
Liquidation value				200.0
Capital gains (losses)				(83.0)
Capital gains taxes				(17.4)
After-tax liquidation value				217.4

Figure 5.6: Dell Tablet - Depreciation Schedule

However, Dell may have to pay taxes when it sells the facility if sells the facility for a price greater than the book value or **tax basis** of the facility. The book value is the original investment value less accumulated depreciation, or $\$400 - 3 \times 39 = \283 million. When the sale price of an asset is greater than the tax basis, a **capital gain** occurs and taxes are owed on the gain. When the sale price of an asset is less than the tax basis, a **capital loss** occurs and a tax shield - reduction in taxes - is realized. Taxes on capital gains/losses are computed as follows.

$$\text{Tax on capital gains} = \underbrace{(\text{Sale Price} - \text{Book value})}_{\text{Capital gains(losses)}} \times \text{Tax rate}$$

In our example, capital gains taxes are

$$(200 - 283) \times 0.21 = -\$17.4 \text{ million.}$$

Because the facility is sold for less than its tax basis, Dell will experience a capital loss and receive a tax shield of \$17.4 million.

The *after-tax* cash inflow from the sale of the production facility in three years time is therefore $200 - (-17.43) = \$217.43$ million. There is a deeper point to be made in recognizing the future cash inflow from the sale of the plant: Assets (and liabilities) don't just disappear. Dell doesn't have to sell the production facility to recognize value in the last year of the project. It could rent or lease the facility. It could re-purpose the facility for

another use. It could continue producing tablets in which case we need to look beyond three years. The point is that when there is still value in an asset, i.e., the asset will continue to generate cash flows, that value should be recognized or we risk understating the value of our project.

5.3.4 Net Working Capital Investment

The last piece of the free cash flow puzzle (equation 5.3) is presented in Figure 5.7, which presents a working capital schedule. This schedule translates the forecast drivers for accounts receivable, inventory, and accounts payable into dollar values.

Working capital schedule (\$mil)	0	1	2	3
Accounts receivable	0.0	2.9	42.4	168.6
Inventory	3.8	48.3	179.0	0.0
Current assets	3.8	51.2	221.4	168.6
Accounts payable	0.0	1.9	24.1	89.5
Current liabilities	0.0	1.9	24.1	89.5
Net working capital	3.8	49.3	197.3	79.1
Recovered net working capital				(79.1)
Net working capital investment	3.8	45.6	148.0	(197.3)

Figure 5.7: Dell Tablet - Working Capital Schedule

We can translate days receivable forecasts into accounts receivable forecasts using the following relation.

$$\text{Accounts receivable} = \underbrace{\frac{\text{Days receivable}}{360}}_{\% \text{ Sales Outstanding}} \times \text{Sales}. \quad (5.5)$$

For example, at the end of the first year of sales, our accounts receivable or what we're owed by our customers is

$$\frac{14}{360} \times 75 = \$2.9 \text{ million.}$$

There is a nice intuition to this calculation. Dell takes 14 days to collect cash on sales to its customers. If sales are distributed uniformly throughout the year, then at any point in time $14/360=3.9\%$ of their sales for the previous 360 days is uncollected. Therefore, at the end of each year, the money owed to Dell by its customers is 3.9% of the annual sales. A quick

aside, 360 days in a year is often used to ensure equal length years and 90-day quarters. That said, using 365, 365.25, or precise day counts are perfectly fine and shouldn't have a noticeable effect on our cash flows or decision criteria.

Sales being “distributed uniformly throughout the year,” is an important assumption for this calculation to make sense. If Sales are seasonal, equation 5.5 can produce poor estimates of receivables at any point in time. Imagine that all tablet sales occur on Christmas. At the end of the calendar year, the true accounts receivable should be the entire sales for the year, 100%. Equation 5.5 will estimate that only 3.9% of the sales are outstanding. While extreme, this example is illustrative. Many companies have seasonal sales, which imply that sales do not arrive uniformly throughout the year. One way to address this issue is to work on a quarterly basis. That is, make each column a quarter in duration, instead of a year.

The inventory policy for Dell is to pre-build 25% of their anticipated sales in the *following* year. Their first year sales are expected to be 0.2 million units at a cost of \$75 per unit. The inventory figure today is therefore $0.25 \times 0.2 \times 75 = \3.75 million. This means Dell starts year one with 0.2 million units in their inventory, costing them \$3.75 million. As time progresses, Dell sells this inventory and at the same time builds more inventory to (i) meet the demand for the rest of the year and (ii) create additional inventory for the start of year two.

Finally, accounts payable is computed in a similar manner to accounts receivable. At the end of each year, Dell will owe its suppliers money. If their costs are distributed evenly throughout the year, we can approximate how much they owe with the ratio of days payable to days in the year, $45/360 = 0.125$. For example, after the first year, Dell should owe its suppliers 12.5% of its \$15 million in COGS. So, an additional \$1.875 million will show up on their balance sheet because of this project.

Three comments on payables are in order. First, computing accounts payable in this manner assumes that purchases are uniformly distributed throughout the year. If, like sales, Dell's purchases are highly cyclical, this can lead to significant error in payables estimates. Second, we are assuming that Dell pays cash for all of its other expenses, overhead and R&D. This need not be true. Likewise, we are assuming Dell is paying cash for its initial inventory today because the accounts payable today, period 0, is zero. This is almost surely untrue, but highlights the difference between cash and credit purchases. Because we assume that Dell pays cash today to acquire its inventory, there is an extra outflow of cash today. If they could postpone that cost by delaying payment using credit, they could increase the value of the project. Pay later is preferred to pay now, all else equal, because of the time value of money.

Net working capital is current assets less current liabilities. Net working capital *investment* is the year-on-year change in net working capital. For year 0, today, NWCI is $3.8 - 0 = \$3.8$ million. This is the net working capital in year 0 minus the net working capital in year “-1”, one year before the start of the project. Because there was no net working before the start of the project, we assign zero to this value. What this means for Dell is that it has to come up with \$3.8 million today to invest in net working capital. Specifically, it has to buy \$3.8 million of inventory. Similarly, Dell has to invest $49.3 - 3.8 = \$45.6$ million in net working capital in year 1. This investment is to build up its inventory and provide credit for its customers less the credit it receives from its suppliers.

The first three years of this project, net working capital is increasing, representing investments that Dell must make - cash outflows. Only in year 3 does net working capital decline, representing a cash inflow for Dell. Reminder: It is important to distinguish between net working capital, which is a number at a point a time (a so-called **stock variable**) and net working capital investment, which is a number corresponding to a period of time (a so-called **flow variable**), like a year. It is the *change* in net working capital that corresponds to the actual investment and cash flow.

At the end of the project in year 3, we show recovered net working capital of \$79.1 million. This figure represents the collection and payment of receivables and payables still outstanding in the last year. More precisely, Dell will collect the \$168.6 million from their customers during the first two weeks of the fourth year. Likewise, Dell will pay their suppliers \$89.5 million during the first 45 days of the fourth year. However, rather than create another column in our spreadsheet, we’ve just assumed that these collections and payments all occur at the end of the last project year. This assumption violates our time value of money rule - never add cash flows arriving or going at different points in time. However, because the cash flows are reasonably far in the future and the time difference small, most financial analysts will overlook this sin.

5.3.5 Free Cash Flows and Valuation

Figure 5.8 puts everything together to compute the project’s free cash flows.

The first two years of negative cash flows correspond to the investment costs of the project. However, the outflows flip to inflows in the second year and take off in year 3. It’s interesting to note the large differences between profits and cash flows. For example, the project loses \$40 million today but requires \$443 million in cash because of the capital

Free cash flow (\$mil)	0	1	2	3
NOPAT	(39.5)	(29.2)	148.4	756.3
Depreciation and amortization	0.0	39.0	39.0	39.0
NLTI	(400.0)			217.4
NWCI	(3.8)	(45.6)	(148.0)	197.3
Unlevered free cash flow	(443.3)	(35.8)	39.5	1,210.0

Figure 5.8: Dell Tablet - Free Cash Flows

expenditures and investment in working capital. Likewise, year 2 profits are more than three times larger than free cash flow because of the continuing working capital investment.

With the free cash flows and the 12% cost of capital, we can compute our decision metrics shown in Figure 5.9

Valuation	0	1	2	3
NPV	\$417.5			
IRR	39.2%			
Cumulative free cash flows	(443.25)	(479.05)	(439.58)	770.38
Payback period (years)	2.4			

Figure 5.9: Dell Tablet - Decision Criteria

The NPV of the project,

$$NPV_0 = -443.3 - \frac{35.8}{(1 + 0.12)} + \frac{39.5}{(1 + 0.12)^2} + \frac{1,210.0}{(1 + 0.12)^3} = \$417.5 \text{ million,}$$

is positive, suggesting we accept the project. The IRR is

$$NPV_0 = 0 = -443.3 - \frac{35.8}{(1 + IRR)} + \frac{39.5}{(1 + IRR)^2} + \frac{1,210.0}{(1 + IRR)^3} \implies IRR = 39.2\%,$$

and is greater than the 12% cost of capital, also suggesting we accept the project. The payback period is just under two and a half years. Depending on the cutoff, the project could be viewed as paying off sufficiently quickly or taking too long to recover the investment.

5.3.6 Comments

The only thing we can be sure of in the preceding analysis is that every single number will be wrong, which begs the question: Why bother? If we know the numbers are wrong, won't

that lead us to an incorrect decision? The answer is an emphatic no! The goal with DCF analysis is not to get the *correct* number. That will never happen because it requires a precise knowledge of the future. One rationale for performing a DCF analysis is to force us to clearly articulate how our ideas will create value. If we can't show, based on our assumptions, that our ideas create value, then we should reconsider them.

Thoguh, we don't always have to, or even want to, perform a full blown DCF analysis for every project that comes across our desk. The most valuable asset we have is time, and there is no rationale to waste it on projects that are unlikely to move forward. For example, if we are working in a product space with operating margins that are typically 25%, and we have a project with a 5% margin, it's going to be an uphill battle to convince leaders that the project makes sense. Likewise, if our boss is expecting revenue growth of 30%, and we propose a project with declining revenues, this project will be difficult to push through.

I'm not saying that decisions should be made based solely on **key performance indicators (KPIs)** like revenue growth and margins, only that they are often good indicators and the focus of many managers. Large differences between these KPIs and management's expectations of them need to be reconciled if arguments predicated on NPV or IRR are to be compelling.

Related, for many managers, working capital and long-term investments are outside the scope of their duties. Their focus is on a P&L. These leaders can still construct pro forma P&Ls to understand the revenue and expense implications of their ideas. The knowledge that others in the company, e.g., operations and finance, may have to deal with other factors related to their ideas provides perspective on why they may receive push back from others in the organization. Understanding the whole picture allows us to engage with other leaders in a productive manner.

5.4 Sensitivity Analysis

Coming up with a baseline valuation as we did above is half the battle. The other half is **sensitivity analysis**. In other words, with the model built, now is the time to start asking "what if?" There are an infinite number of scenarios we could examine. We'll explore a few here and introduce some useful metrics and analyses.

A couple of comments before we set off. First, with an appropriately designed financial model - see the accompanying Excel spreadsheet - answering what-if questions requires nothing more than changing a number in a cell and seeing what happens to NPV, IRR, and

other key performance metrics (margins, earnings, etc.). Second, we need to understand the value of this analysis. Imagine someone asks: “What happens if market demand is lower than we expect?” We don’t need a financial model to tell us that the project KPIs and value are going to look worse. What we do need a financial model for is understanding *how much* worse. There is a big difference between value declining but remaining positive, and value turning negative. Our model can help distinguish between these scenarios and in doing so quantify our **risk exposure**.

5.4.1 Market demand and pricing

- *What if market growth is slower than expected so that forecasted tablet demand is half that of our baseline estimates?* Table 6 shows the implications of unit sales falling to 0.5, 5.0, and 15.0 million from 2013 to 2015. NPV falls from \$417.5 to \$56.2 million;

Key Performance Indicator	Baseline	Slow growth
NPV	417.5	56.2
IRR (%)	39.2	16.5
Payback period (years)	2.4	2.4
1st year NOPAT	-29.2	-35.9
1st year FCF	-35.8	-19.6

Table 6: KPI Response to Lower Market Sales (\$mil)

IRR falls from 39.2% to 16.49%. From a value perspective, this demand shock isn’t a great risk. The project is still value-accretive. The payback period gets pushed out but only modestly - a little over two months. Year one earnings losses jump from \$29.2 to \$36.0 million, which will need to be explained to investors. Interestingly, liquidity pressure decreases when market demand is lower because less investment in net working capital is needed - fewer sales means fewer accounts receivable, which spike in the first year. Year one free cash flow increases from -\$35.8 to -\$19.6 million.

- *Price-quantity pairs.* Figure 5.10 presents a two-way data table in which each row corresponds to a different initial unit price, and each column corresponds to a different market share. The 10% annual price growth is held constant. The numbers inside the table correspond to the project NPVs for the price-quantity pairs. For example, our baseline assumptions of \$400 unit price and 20% market share generate a \$417.5 million NPV, identified by the black outlined cell.

The color formatting indicates the progression of NPV from negative (red) to baseline (yellow) to maximum (green). While everyone would want to be in the lower right

Unit Price (\$)	Market Share						
	10%	15%	20%	25%	30%	35%	40%
300	(42.3)	28.1	98.6	169.1	239.5	310.0	380.5
320	(0.9)	80.7	162.4	244.0	325.7	407.3	489.0
340	40.5	133.3	226.2	319.0	411.8	504.6	597.4
360	81.9	185.9	289.9	393.9	497.9	601.9	705.9
380	123.3	238.5	353.7	468.9	584.1	699.2	814.4
400	164.8	291.1	417.5	543.8	670.2	796.5	922.9
420	206.2	343.7	481.3	618.8	756.3	893.8	1,031.4
440	247.6	396.3	545.0	693.7	842.4	991.2	1,139.9
460	289.0	448.9	608.8	768.7	928.6	1,088.5	1,248.3
480	330.4	501.5	672.6	843.6	1,014.7	1,185.8	1,356.8
500	371.9	554.1	736.3	918.6	1,100.8	1,283.1	1,465.3

Figure 5.10: Dell Tablet NPVs for Different Price-Quantity Combinations

corner of the matrix, at the highest NPV, consumer demand doesn't work that way. Typically, demand curves slope down, implying quantity demanded declines as price increases. How fast depends in large part on the slope of the demand curve, a topic for an economics course. The bottom line is that someone in the company - e.g., sales or marketing - needs to understand by how much demand will fall when prices increase to identify the optimal price quantity combination that maximizes NPV. The matrix is simply a useful visual to lay out different possibilities.

5.4.2 Cost Assumptions

- *What if inflation leads to faster unit cost increases of 7% per year as opposed to 3%?* Again, the qualitative implications of an increase in input costs are obvious; value

Key Performance Indicator	Baseline	7% Inflation
NPV	417.5	379.6
IRR (%)	39.2	36.9
Payback period (years)	2.4	2.4
Average operating leverage	3.7	3.5
1st year NOPAT	-29.2	-29.2
1st year FCF	-35.8	-37.7

will decline all else equal. By how much is less clear. The table shows that project value is only modestly affected. Since the cost increases only hit in years 2 and 3 of the project, there is no impact on first year NOPAT, though first year free cash flow is reduced because inventory costs, based on next year's sales, increase. Average operating leverage - the ratio of fixed costs-to-variable costs - for the project decreases

because variable costs increase relative to overhead. The operating risk of the project decreases.

- *R&D investment will be instrumental to future product iterations in a fast-moving product life cycle. What is the impact of an increase in the R&D budget to \$25 million per year after the initial investment?* Value decreases but only modestly. First year earn-

Key Performance Indicator	Baseline	7% Inflation
NPV	417.5	389.0
IRR (%)	39.2	37.3
Payback period (years)	2.4	2.4
Average operating leverage	3.7	4.0
1st year NOPAT	-29.2	-41.1
1st year FCF	-35.8	-47.7

ings and cash flow is significantly negatively impacted. Operating leverage increases because of the additional overhead.

5.4.3 Key Value Drivers and Break Even Points

- *Which assumptions are most important for the success of the project?* In other words, to which assumptions is the project's value most sensitive, i.e., the key value drivers? Table 7 presents the elasticity - and intermediate calculations - for some of the project's assumptions. **Elasticity** measures the percentage change in a variable for a one percent change in another variable. In our setting, elasticity measures the percentage change in the value of the project (NPV) for a 1% increase in the forecast driver. For example, the elasticity of project value with respect to the unit price is computed as follows.

$$\frac{\text{Change in NPV}}{\text{Change in unit price}} \times \frac{\text{Unit price}}{\text{NPV}} = \frac{430.2 - 417.5}{404 - 400} \times \frac{400}{417.5} = 3.06$$

Comparing elasticities shows that the key value drivers are the unit sales price and SG&A, with the largest magnitude elasticities of 3.06 and -4.60. A 1% increase in the unit sales prices leads to a 3.06% increase in NPV, whereas a 1% increase in SG&A leads to a 4.6% *decrease* in NPV (note the negative sign). Assumptions regarding days receivable and payable, for example, are largely irrelevant for the valuation based on their elasticities, -0.01 and 0.01 respectively.

A word of warning: Elasticities are calculated assuming all other parameters are unchanged. In practice, when we change one assumption we will also be changing one or more other assumptions perhaps intentionally. For example, changing the sales

Forecast driver	Baseline	Baseline + 1%	Elasticity (%)
Dell market share (%)	20.0	20.2	1.21
NPV	417.5	422.5	
Unit price (\$)	400.0	404.0	3.06
NPV	417.5	430.2	
Loss of laptop revenue (\$mil)	5.0	5.1	-0.05
NPV	417.5	417.3	
Unit cost (\$)	75.0	75.8	-1.32
NPV	417.5	411.9	
Unit cost growth (%)	10.0	10.1	-0.28
NPV	350.3	349.3	
SG&A (% of sales)	60.0	60.6	-4.60
NPV	417.5	398.3	
R&D (\$mil)	10.0	10.1	-0.05
NPV	417.5	417.3	
Capital expenditures (\$mil)	400.0	404.0	-0.81
NPV	417.5	414.1	
Days receivable (days)	14.0	14.1	-0.01
NPV	417.5	417.4	
Inventory (% next year output)	25.0	25.3	-0.05
NPV	417.5	417.3	
Days payable (days)	45.0	45.5	0.01
NPV	417.5	417.5	
Tax rate (%)	21.0	21.2	-0.35
NPV	417.5	416.0	
Cost of capital (%)	12.0	12.1	-0.67
NPV	417.5	414.7	

Table 7: Dell Table: Key Value Drivers

price of the unit is almost surely going to change the quantity sold. This is the law of demand. Nonetheless, this exercise is particularly useful for understanding where we should spend most of our time and attention. Specifically, we need to invest in understanding and getting as accurate an assumption as possible for our key value drivers (high elasticity), and less time and attention on the relatively unimportant value drivers (low elasticity). Just make sure to remember that “high” elasticities are those that are large *in magnitude*, i.e., regardless of sign.

- *At what point does the project become value-destructive for each of our assumptions?*

Table 8 presents the baseline and break-even values for each forecast driver. The

Forecast driver	Baseline	Break-even
Dell market share (%)	20.0	3.5
Unit price (\$)	400.0	269.1
Loss of laptop revenue (\$mil)	5.0	104.3
Unit cost (\$)	75.0	131.6
Unit cost growth (%)	3.0	41.5
SG&A (% of sales)	60.0	73.1
R&D (\$mil)	10.0	230.0
Capital expenditures (\$mil)	400.0	894.1
Days receivable (days)	14.0	1,512.8
Inventory (% next year output)	25.0	538.8
Days payable (days)	45.0	(8,336.9)
Tax rate (%)	21.0	81.6
Cost of capital (%)	12.0	39.2

Table 8: Dell Table: Baseline and Break-Even Values

break-even values must be calculated numerically. We can use GoalSeek in Excel. (See the appendix C for details on this function. For example, at a unit cost \$131.60, the NPV of the project is zero. Anything anything than \$131.60 and the the NPV will be negative and the project value destructive. Likewise, SG&A cannot go above 73.1% of sales or else the project will lose money.

For each, forecast driver, we have to ask: Are values beyond the break-even value (i.e., values leading to negative NPV) plausible and, if so, how likely? We have limited control over units costs, which vary largely because of fluctuations in labor and materials costs. Can events lead to an increase above \$131.6? And, if they can, how much higher can they get? The goal is to understand our risk exposure: how likely we are to lose money and how much we can lose.

For example, assuming we can keep cost overruns on plant development under \$494.1 million (\$894.1 - \$400), the project will still be value-accretive (NPV positive).

The days receivable and payables break-even values convey a message similar to that conveyed by their small elasticities. As long as we can collect from customers within 1,512.8 days (a little over four years), the project NPV will be positive. This is clearly not a concern and indicative of the small impact days receivables has on the project's NPV. Likewise, as long as our suppliers don't make us pay 8,336.9 days (22.8 years) *in advance* of our purchases, the project NPV will be positive. This number is obviously absurd suggesting that days payable is not a risk factor with which we should be concerned. Though, this is not to say that there is no risk from suppliers (supply chain problems, tariffs, materials shortages, etc.).

5.4.4 Targeting KPIs

- *How can the project forecast drivers be modified to achieve a target IRR of 45%?* Table 9 presents the baseline and target values for a subset of forecast drivers that achieve the target IRR of 45%. Setting targets for employees is a common tool used by leaders.

	Target KPI	Baseline	Target driver
IRR (%)	45.0		
Market share (%)		20.0	24.5
Unit cost (\$)		75.0	62.0
Capital expenditures (\$mil)		400.0	343.2

Table 9: Dell Tablet: Targeting the IRR

This analysis can determine whether a project can achieve those targets. In the current example, we can ask whether it is plausible to achieve a 24.5% initial market penetration, or whether we can squeeze our suppliers on their prices, or if we can negotiate or pare down the up-front investment. If the answer is no, then achieving the target 45% IRR may be infeasible with this project. Of course, this analysis doesn't answer the question of why an IRR of 45% is a goal for which a company should strive.

- *How much of the tablet project's cash flow would Dell be willing to spend lobbying Congress to lower the corporate tax rate to 10%?* At a 10% corporate tax rate, the project NPV is \$493.3 million, a \$75.8 million increase. So, the most Dell would be willing to spend is \$75.8 million, assuming the lobbying process guaranteed the tax reduction. (Good luck with that.) If instead the lobbying process has a 20% chance of

success, then Dell might be willing to pay up to the expected benefit from lobbying, which is $0.20 \times \$75.8 = \15.16 million.

5.4.5 Scenario Analysis

The sensitivity analysis performed thus far changes one, at most two, assumptions at a time. But, what if we want to explore the effects of changing more than two assumptions? This is what scenario analysis accomplishes.

- *What is the impact of a recession on the tablet project assuming the following changes. Tablet demand would be cut in half each year. Market share would increase to 25% as consumers turn to Dell's lower priced offering. Unit price would be cut to \$350. Cost per unit would decline from \$75 to \$70. Changing all of the relevant parameters shows that NPV falls to \$47.4 million, still positive but only marginally so.*

To summarize, sensitivity analysis is a critical part of any discounted cash flow analysis. It can provide confidence in our decisions, address concerns about our assumptions, and quantify our risk exposure.

5.5 Key Ideas

Corporate financial decision making is fundamentally no different from personal decision making. Companies, like us, want to undertake decisions with positive net present value. In a corporate setting, the cash flows used to measure the costs and benefits are called free cash flows and are measured in a specific way (equation 5.3). The discount rate, r , is referred to as a cost of capital. Despite the new lingo, the mechanics and intuition are identical to personal financial decisions. We lay out the cash flows on a timeline, discount them, and then sum to estimate value or net value.

- There are three primary decision criteria used for assessing stand alone projects: NPV, IRR, and payback period.
 - NPV should always be used for decision making.
 - IRR is a useful communication device but has significant shortcomings, especially when comparing projects, that can lead decision-makers astray.

- Payback period is informative about the risk and liquidity needs of a project but should not be relied upon by itself for decision making.
- Free cash flow measures the financial benefits net of costs in corporate settings.
- The only way a decision affects value is through cash flows or discount rates, the latter of which is dictated primarily by financial considerations we'll discuss later. Therefore, a clearly articulated business plan explicitly links strategic and operating decisions to the value channels contained in free cash flow - sales, expenses, taxes, and investment (long-term and short-term).
- A critical element of any business decision is sensitivity analysis, which explicitly recognizes the limitations of financial analysis and helps quantify risk exposure.

5.6 Technical Appendix

Linear interpolation is sometimes used to estimate the payback period when if (i) falls between periods and (ii) the cash flows are uniformly distributed throughout the period. Intuitively, linear interpolation estimates values between two points by assuming all values between the two points lie on a line. Figure 5.11 illustrates this intuition with the payback period for the Dell tablet project. The horizontal axis are the cumulative free cash flows; the vertical axis is the time period.

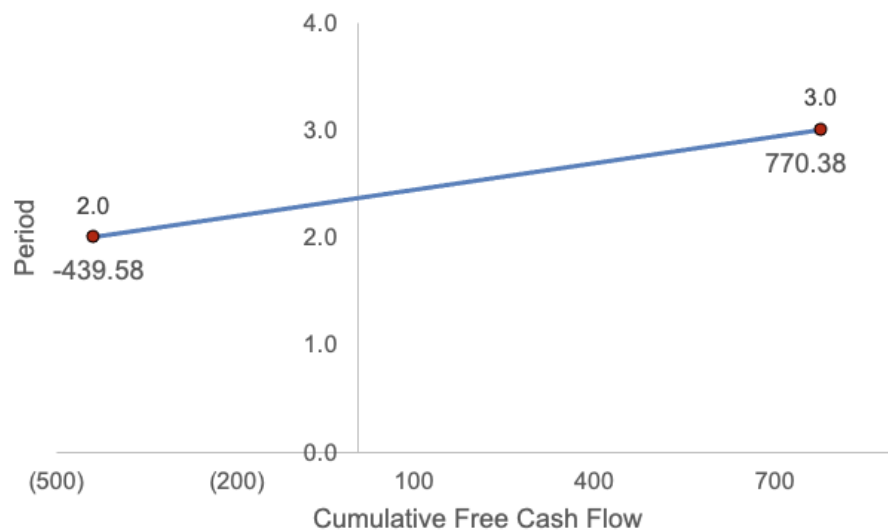


Figure 5.11: Linear Interpolation of Cumulative Free Cash Flows from Dell Tablet Project

The left and right points in the figure correspond to the cumulative free cash flow at the end of the second and third years, respectively. We need to find where the line crosses the vertical axis, i.e., where the cumulative free cash flow exactly equal 0. This is our linearly interpolated payback period. Recall that with two points, we can identify the line connecting them. The line in which we're interested is

$$\text{Period} = \text{Slope} \times \text{Cumulative Free Cash Flow} + \text{Intercept}$$

The slope of the line is “rise over run,” or

$$\text{Slope} = \frac{3.0 - 2.0}{770.38 - (-439.58)} = 0.0008264736024$$

The intercept can be found using any point on the line in conjunction with the slope.

$$\begin{aligned} \text{Intercept} &= \text{Period} - \text{Slope} \times \text{Cumulative Free Cash Flow} \\ &= 3 - 0.0008264736024 \times 770.38 = 2.3633012662 \end{aligned}$$

Thus, the line connecting these two points is

$$\text{Period} = 0.0008264736024 \times \text{Cumulative Free Cash Flow} + 2.3633012662.$$

Because we are interested in the period at which the cumulative free cash flow equals zero, the intercept of our line gives us the linearly interpolated payback period. More generally, given two points (x_0, y_0) and (x_1, y_1) , the linearly interpolated value for y at the point x is given by the following expression.

$$y = y_0 + (x - x_0) \frac{y_1 - y_0}{x_1 - x_0}$$

5.7 Problems

For all problems requiring calculation, it is strongly recommended - in many cases required - that a spreadsheet or other computing program be used.

5.1 (*Revenue and expense attribution, incremental earnings*) Frito-lay is introducing its new Dorito-brand chip called “Fireball,” a super spicy corn chip. It is expecting to sell one million bags next year, a figure that will grow by 5% per year thereafter. Each bag of chips sells for \$1.99 and there is no plan to change that price.

The introduction of Fireball chips will lead a number of customers to switch from Nacho cheese flavored chips. Specifically, 100,000 fewer bags of Nacho cheese chips will be sold when Fireball is introduced, a number that is expected to grow by 3% per year. Nacho cheese chips sell for the same price as Fireball chips.

Fireball and Nacho cheese production costs are \$0.47 per bag (e.g., packaging, labor, and ingredients). Marketing expenses are \$500,000 during the first year of sales and \$200,000 each year thereafter. Because of the anticipated reduction in Nacho Cheese chip sales, marketing spend on Nacho cheese chips will be limited to certain geographic regions, resulting in a \$100,000 per year reduction in Nacho Cheese marketing expenses.

Using this information, answer the following questions.

- a. What are the incremental sales for the Fireball chip over the next 10 years?
- b. How would your answer to the previous question change if half of the reduction in Nacho Cheese sales was due to consumers who were just tired of Nacho cheese and would have stopped buying Dorito chips altogether were it not for the introduction of Fireball?
- c. What are the incremental expenses for the Fireball chip over the next 10 years.
- d. What are the incremental operating earnings (EBITDA) and operating margins for Fireball over the next 10 years? What is the compounded annual growth rate (CAGR) of EBITDA over the 10 years?

5.2 (*Revenue attribution, incremental sales*) Ping is a golf club and accessory manufacturing company that manufactures and sells golf equipment and apparel. They are preparing to introduce a new golf club, the Ping Eye 42, which has projected first-year sales of \$10,000,000. In addition, Ping estimates an increase in sales (i.e., “lift”) of balls and apparel equal to \$1,500,000 as a result of the new club release. What are the first-year incremental sales of the Ping Eye 42?

5.3 (*Free cash flow*) Max has a corner lemonade stand. From 3:00 PM to 5:00 PM on Tuesday, Max and his sister, Sophie who works for him, sold 27 cups of lemonade at a price of \$1.50 per cup. Each cup of lemonade costs \$0.25 in materials - cups, lemons, sugar. In addition to his direct costs, Max pays his sister an hourly wage of \$2.00 per hour. He also pays “taxes” to his father equal to 10% of his profits. Max keeps any money left over.

Using this information, answer the following questions.

- a. What were Max’s daily
 - i. sales?
 - ii. expenses?
 - iii. operating earnings (EBIT)?
 - iv. NOPAT?
 - v. free cash flow?

5.4 (*Recovered net working capital*) John’s microphone project is ending today. While perusing the project balance sheet, John notices the following current accounts.

Account	Amount (\$)
Inventory	2,000
Accounts receivable	600
Accounts payable	900

What happens to these accounts now that the project is ending? What are the implications for free cash flow? (Hint: Think about the business.)

5.5 (*Sale, receivables, free cash flow*) Billy-Jean jeans is a wholesaler of designer jeans. Their most recent and projected annual jean sales are as follows.

	December of Year				
	2021	2022	2023	2024	2025
Sales (\$000s)	500	800	1,000	1,300	1,800

Billy-Jean’s collection policy is 90 days, meaning they allow their customers to pay 90 days after their purchase.

Using this information, answer the following questions/perform the following tasks.

- a. Compute the year-on-year sales growth rates implied by Billy-Jean’s forecasts.

- b. What is the compounded annual growth rate (CAGR) of Billy-Jean's sales over the four years?
- c. Construct annual forecasts of Billy-Jean's end-of-year accounts receivables assuming that sales are evenly distributed throughout the year. (I.e., They sell the same amount each day, month, quarter, etc.) What are the free cash flow implications of Billy-Jean's receivables policy under this sales pattern?
- d. Reconstruct annual forecasts of Billy-Jean's end-of-year accounts receivables assuming that all sales occur at the *end* of the year during holiday season. What are the free cash flow implications of Billy-Jean's receivables policy under this sales pattern?
- e. Reconstruct annual forecasts of Billy-Jean's end-of-year accounts receivables assuming that all sales occur at the end of school, in May and June. What are the free cash flow implications of Billy-Jean's receivables policy under this sales pattern?
- f. How can we get an accurate assessment of the cash flow implications of accounts receivable when sales are highly seasonal?

5.6 (*Free cash flow, inventory financing, accounts payable*) Coltrane Pet is a small manufacturer of pet supplies. A recent audit of their recent financial statements revealed no accounts payable on their balance sheets, but quite a bit of inventory. How is this possible? What are the implications of this zero payables strategy Coltrane appears to be following?

5.7 (*Sunk costs*) RennerTech is small bio-tech firm located in La Jolla, California. Their mRNA therapeutic, Rexall, has just cleared Phase 3 of the FDA trials. Unfortunately, a competitor's drug with similar indications has also cleared Phase 3 of the FDA trials. RennerTech is trying to determine whether they should go to market with Rexall by performing a discounted cash flow analysis.

One sticking point in the analysis is how to account for the costs of the FDA trials, which were substantial. The CEO, Jeremy Strothers, argues: "We spent over \$50 million on the FDA trials, we can't just ignore that in our DCF analysis." Is Jeremy correct? Should the money RennerTech spent on moving through the FDA trails be accounted for in their DCF analysis? Explain your answer.

5.8 (*Expense allocation*) Tim Ferris is the head of financial planning and accounting (FP&A) at Axiom Inc., a telecommunications company. Tim has received word of a new project

that would expand internet capacity in the Northwest United States. In response, he has decided to allocate several million dollars of existing administrative expenses to the new project. How exactly would this expense affect the free cash flow of the Northwest expansion project?

5.9 (*Free cash flow considerations*) Ringchange LLC is assessing a new tungsten mining project at their mine in the Gobi desert. Currently, Ringchange is extracting manganese from the site. However, with additional drilling and new technology, they will be able to extract the higher priced tungsten instead of manganese. Which of the following considerations should be accounted for in the calculation of free cash flow for the project?

- a. The costs of the geological surveys performed the previous year.
- b. The effects of reduced competition in tungsten markets.
- c. The depreciation of new mining equipment.
- d. The depreciation of existing mining equipment.
- e. The clean up costs associated with closing the mine when fully depleted.
- f. The CEO's salary.
- g. The company's tax savings arising from forecast initial losses at the mine.
- h. The interest expense on the debt used to finance the purchase of the mine site.
- i. The lost after-tax profits from Ringchange's manganese operation.

5.10 (*Free cash flow, asset liquidation*) Salvagers is an aluminum recycling company. Their sorting machine is at the end of its life and about to be replaced. The machine was purchased eight years ago for \$12 million. It has been depreciated in a straight-line manner over its eight year life to its current salvage value of \$2 million. The company is planning on selling it to a smaller competitor for \$3.4 million. If Salvagers marginal tax rate is 18%, what are the cash flow implications of this transaction for Salvagers?

5.11 (*Depreciation tax shield, accelerated depreciation*) Modified Accelerated Cost Recovery System (MACRS) is an asset depreciation system. Unlike straight-line depreciation, MACRS front-loads depreciation expenses; hence the modifier accelerated. For example, office furniture has a depreciable life of seven years according to the IRS. A desk that costs \$1,000 would have the following depreciation schedules.

(MACRS uses a "half-year" convention that treats assets as though they were put in use halfway through the year. As a result, the IRS allows for an extra half-year of

	Years							
	1	2	3	4	5	6	7	8
Straight-line (\$)	143	143	143	143	143	143	143	
MACRS (\$)	143	245	175	125	89	89	89	45

depreciation before the asset is sold or retired.) The different depreciation schedules are illustrated in Figure 5.12.




Figure 5.12: Straight Line versus MACRS Depreciation

Using the desk example, compute the depreciation tax shield for both straight line and MACRS depreciation assuming a cost of capital of 10%, and a tax rate of 21%. Which depreciation approach - straight line or MACRS - provides more value to firms? Is there a reason the firm might choose not to employ that value maximizing approach?

5.12 (*DCF, decision criteria*) For a project with the following cash flows, calculate the

- Net present value,
- Internal rate of return,
- Payback period,
- Discounted payback period, and
- Multiple on invested capital or MOIC (computed as the ratio of total cash inflows to total cash outflows without regard for the time value of money)

	Period			
	0	1	2	3
Cash flows	-1000	200	400	800

Assume that the appropriate cost of capital is 8.5%.

5.13 (*DCF, decision criteria, dynamic valuation*) Max had so much fun selling lemonade, he decided to make a business of it. However, to get financing from his parents, he had to value his business. Max would run the operation starting tomorrow for five days, after which school started. The stand would be open for two hours per day and his sister agreed to work every day. His forecasts for the business were as follows.

- Daily cups sold would start at 30 on the first day of business, and grow by 10% per day thereafter because of word of mouth.
- His retail price per cup would stay constant at \$1.50.
- His cost per cup would stay constant at \$0.25.
- His sister's hourly wage would stay constant at \$2.00.
- At the end of each day beginning today, he would have to hold \$3 of lemons, sugar, and cups in inventory. On the last day, he would sell out all of his inventory.
- Because Max operated an entirely cash business - cash payments from customers and cash payments to his suppliers (mom and dad), he did not anticipate any receivables or payables.
- He would need to build a proper lemonade stand today at a cost of \$12. The stand would depreciate in value equally over the five days to a salvage value of zero.
- Because of his initial success, Dad decided to tax Max 15% on his daily profits.
- Max's daily cost of capital is 1%. (It's a high risk business.)

Using this information, answer the following questions.

- What is the pro forma P&L statement for the lemonade business? Specifically, what are the revenues, expenses, and (before- and after-tax) profits?
- What are the free cash flows for the business?
- What is the NPV, IRR, and payback period of Max's business?

- d. *Advanced* After watching Max's success over the first two days, his sister decides to buy the business at the end of the second day? What is the business worth? I.e., what is the value of the business as of the end of the second day? (Hint: The value of an asset is the sum of the *future* discounted cash flows generated by the asset.)

5.14 (*DCF, valuation, break even analysis*) Your brother has asked you to invest in his new company. He is asking for \$10,000 today and promising to give you back \$20,000 in four years. Knowing the business - and your brother - you estimate that the opportunity cost of this investment is 30% per annum.

Using this information, answer the following questions.

- What are the NPV, IRR, and Payback period for this investment opportunity?
- Based on your answers to the previous question, would you invest in this opportunity? Explain your answer.
- What would your brother have to pay you in four years for you to break even?
- What opportunity cost of capital would the investment need to break even?

5.15 (*DCF, valuation, break even analysis*) Substack is an online platform enabling content producers - e.g., writers, bloggers, podcasters - to sell their content. Substack keeps approximately 18% of all revenue to cover their costs and credit card fees. Jerry has convinced all 10 of his family members - immediate and distant - to subscribe to his newsletter, which costs \$7 per year and is payable at the start of each year. Jerry plans on writing his newsletter for 20 years and his family members have promised to continue their subscriptions for as long as he continues to write. Jerry's opportunity cost of capital is 8% per year and his personal tax rate is 32%.

Using this information, answer the following questions.

- How much revenue will Jerry generate each year?
- What are Jerry's before- and after-tax earnings each year?
- What is the value of Jerry's newsletter in today's dollars?
- (*Advanced*) Xtract Inc. is an online marketing company that helps content producers grow their business. They have promised to deliver 20% annual growth in the number of Jerry's subscribers. What is the most Jerry would be willing to pay for Xtract's services assuming payment for 20 years of service was made today?

e. (*Advanced*) If Jerry is allowed to make 20 equal annual payments to Xtract Inc beginning today, what is the most he would be willing to pay each year?

5.16 (*DCF, decision criteria, break even analysis, targeting KPIs*) Hooker Chemical manufactures cleaning solvents and synthetic resins. Hooker is considering building a new plant that will cost \$10 million today. Once online, the plant is expected to generate \$4.5 million of free cash flow one year from today. These cash flows are forecast to grow at 5% per year for the next nine years. At the end of the 10th year, the plant will shutdown.

The manufacturing process generates a significant amount of toxic waste that must be cleaned-up after the plant shuts down. Hooker estimates annual clean-up costs will be \$500,000 per year beginning one year after the plant shuts down and last for 10 years. Hooker Chemical's cost of capital is 12% per annum. However, because the storage and maintenance fee *must* be paid every year, the appropriate discount rate is the risk-free rate of 4% per annum.

Using this information, answer the following questions.

- What are the NPV and IRR of Hooker Chemical's new plant?
- Based on your answers to the previous question, should Hooker invest in the new plant? Explain your answer.
- What is the break-even annual clean-up fee?
- If Hooker is targeting a 20% annual return on investment (IRR), what must the clean-up fee be to achieve this goal?

5.17 (*Valuation and cost of capital relation*) Which is likely more sensitive to changes in the cost of capital: a short-horizon project or a long horizon project? Why? *Advanced:* Compute the following cross partial derivative for the NPV of a project that generates one cash flow at a future time T .

$$\frac{\partial^2}{\partial r \partial T} \left(\frac{CF}{(1+r)^T} \right)$$

What is the sign - positive or negative - of this derivative? What does the sign imply for the impact of changes in the cost of capital on value as the horizon changes?

5.18 (*NPV, IRR, reinvestment return*) One way in which NPV and IRR differ is in the assumed return on any cash flows generated by a project. What rate of return is assumed to be earned on project cash flows when using the NPV criterion? IRR

criterion? Use the following example project to support your answer. The project costs \$100 today and will generate cash flows of \$50 per year over the next three years. The cost of capital for this project is 10% per annum.

5.19 (*DCF, decision criteria*) Frank Dewey Esquire from the law firm of Dewey, Cheatum, and Howe, has been offered an upfront retainer of \$30,000 to provide legal services over the next 12 months to Taggart Transcontinental. In return for this upfront payment, Taggart Transcontinental will have access to 8 hours of legal services from Frank for each of the next 12 months. Frank's billable rate is \$250 per hour and his annual cost of capital is 12%.

- a. What is the NPV of the Taggart agreement? Does it make sense for Frank to accept these terms using the NPV criterion?
- b. What is the IRR of the Taggart agreement? Does it make sense for Frank to accept these terms using the IRR criterion?
- c. How can you reconcile your answers to the previous two questions?

5.20 (*DCF, labor contract valuation*) In 2019, Bryce Harper signed a 13-year contract to play baseball with the Philadelphia Phillies. The contract specified:

- \$20 million signing bonus, and
- an average \$27.5 million per year salary.

Using this information, answer the following questions assuming the signing bonus is paid at the start of 2019 and his salary is paid at the end of each year.

- a. If the Phillie's cost of capital is 12% per year, what is the cost to the Phillies of Bryce's contract in present value terms?
- b. If the Phillies want to buyout the remainder of Bryce's contract five years from today, how much will it cost them assuming his salary is guaranteed and the Phillies' cost of capital is unchanged from 12%?
- c. The Phillies anticipate Bryce's signing will generate an immediate increase in free cash flow of \$40 million to the organization from increased ticket and merchandise revenue. However, this increased cash flow is expected to decline by 4% per year over the life of Bryce's contract. Does his contract make financial sense for the Phillies? Explain your answer.

- d. Bryce prefers to front-load his contract with a first year salary of \$40 million to purchase a home. By what fixed amount must his salary decline each year to maintain the same cost of the contract to the Phillies?
- e. Bryce has a performance clause in his contract that provides him with up to an additional \$800,000 per year should he achieve certain milestones (e.g., becoming an all-star, winning a golden glove or most valuable player (MVP) award). If he has a 25% chance in any year of earning the additional \$800,000 dollars, what is the present value of all the performance costs?

5.21 (*DCF, decision criteria, cash flow decomposition*) The ABB Group is going to invest in a new robotic machine that it hopes will dramatically increase productivity in its engineering division over the next year. The purchase price of the machine is \$850,000, which depreciates an equal amount every three months over the one year project horizon to a salvage value of \$100,000. The firm estimates that it will be able to sell the machine on the secondary market for \$250,000 at the end of the year just after the last depreciation expense is recognized. There is no working capital investment required for this project.

Management estimates that the machine will increase EBITDA by \$150,000 each month during the year beginning at the end of the first month. The company faces a marginal tax rate of 35% and a project cost of capital equal to a 10% APR compounded monthly. You can assume that the company can monetize the tax shield associated with any net operating losses by offsetting taxable income elsewhere in the company.

Using this information, answer the following questions.

- a. What is the monthly periodic cost of capital?
- b. Compute the depreciation expense for each month in the life of the project.
- c. Estimate the after-tax proceeds associated with the sale of the machine at the end of the year.
- d. Estimate the net operating profit after taxes for each month of the project.
- e. Estimate the free cash flows for each month of the project.
- f. What is the net present value (NPV) of the project? According to the NPV criterion, should you purchase the machine and why or why not?
- g. What is the annual internal rate of returns (IRR) on the project? According to the IRR criterion, should you purchase the machine and why or why not?

- h. Compute the net present value of the project by recognizing that the free cash flows can be constructed as the sum of the following components.
 - i. the initial investment,
 - ii. a monthly annuity with cash flows equal to the free cash flows ignoring D&A,
 - iii. a quarterly annuity with cash flows equal to the difference between the free cash flow with D&A and the free cash flow without the D&A, and
 - iv. the difference between the month 12 free cash flow and the free cash flow at the end of a quarter.

5.22 (*DCF, decision criteria, breakeven analysis, salvage values*) Johnny Quaker is an incoming MBA student, who is starting a new business to help pay for his degree. The idea is simple: over the next two years, he is going to sell Wharton gear from a truck parked on Locust Walk. The truck will cost him \$6,000 today, at the start of his first year, and is depreciated in a straight-line manner over a five year usable life to a salvage value of \$500. When he graduates, he estimates that he will be able to sell the truck for \$4,000 when he closes up the business.

He estimates he will need to purchase \$10,000 of inventory at the start of each year, which he plans on selling at a 100% markup (i.e., twice his cost). We can assume that Johnny's is an all-cash business - purchases and sales. Johnny faces a 25% tax on all profits and capital gains, and his cost of capital is 7%.

Using this information, answer the following questions.

- a. How much money does Johnny need to launch the business?
- b. Estimate the NPV, IRR, and payback period for Johnny's venture.
- c. What is Johnny's break-even markup on his product, expressed as a percent?

5.23 (*DCF, decision criteria, buy or lease analysis, break even analysis*) Johnson & Johnson's drug development group is considering the purchase of a new mass spectrometer. The machine has an 8-year life and is estimated to save the company \$4,500 per year in operating costs beginning one year after purchase. The machine would be depreciated on a straight-line basis to a zero salvage value. The company faces a 34% tax rate and a 12% annual cost of capital for the project.

Using this information, answer the following questions.

- a. If the machine costs \$15,000, should it be purchased?

- b. If J&J has the option to lease the machine for \$4,000 per year payable at the end of each year of its eight year life, should they buy the machine or lease it? Assume that the lease payments are expensed each year through an operating lease.
- c. What is the maximum lease payment the J&J would be willing to pay if it was to consider a leasing alternative?

5.24 (*DCF, valuation, target KPIs, contracting business*) At the end of 2021, the Washington Post must determine what to do with its print newspaper business. Circulation has declined from 726,000 to 159,040 readers between 2004 and 2021. A daily print subscription, now only available to readers in the Washington DC area, costs readers \$299. The cost to the Post for printing and delivering the newspaper consists of two components. There is a fixed annual cost of \$15,000,000 corresponding to the support, operation, and maintenance of the printing presses and delivery of the newspapers. Additionally, there is a \$47 per subscriber cost corresponding to materials and delivery costs (e.g., gas). The printing presses and other assets supporting print subscriptions have been fully depreciated so that there is no more depreciation expense after 2021.

As of the end of 2021, The Post forecasts print subscriptions will continue to contract at 9% per year for the foreseeable future. As a result, there will be no further capital investment and net working capital is small enough to be ignored. The Post's cost of capital is 8% per annum and its marginal tax rate is 21%.

Using this information, answer the following questions assuming today is the end of 2021.

- a. What is the cumulative annual growth rate (CAGR) in the Post's print subscriber base from 2004 to 2021?
- b. What are the free cash flows for the print newspaper business from 2022 to 2040? (Hint: You'll need to forecast subscribers, cost per subscriber, fixed cost, etc.)
- c. What is the value of the print newspaper business if the Post continues to run it through 2040?
- d. Should the Post run its print newspaper business through 2040? If not, when should it stop, before or after?
- e. If Congress were to raise the corporate tax rate from 21% to 35% for 2022 and beyond, how would this change affect when the Post should cease print operations and the corresponding value of the business?

- f. Assume inflation and asset obsolescence leads to a “fixed” annual cost that grows at 8% per year starting in 2022. (The cost is fixed in that it doesn’t vary with the number of subscribers.) For example, the fixed annual cost in 2022 is $15,000,000 \times (1 + 0.08) = \$16,200,000$. How would this change affect when the Post should cease print operations and the corresponding value of the business?

5.25 (*DCF, valuation*) Monica Williams is an orthodontist looking to take a three-year hiatus from her practice to spend more time with her young children. Before doing so, she wants to understand the financial consequences. Her projected revenue for this year is \$3.5 million, which is expected to grow at 3% per year thereafter. Her operating profit margin is 60%, and her corporate tax rate is 39%. Monica maintains a 90 days receivable policy, and her current accounts receivables are \$750,000. Her opportunity cost of capital is 14%, and she believes that this decision will have no impact on her practice beyond three years when she returns to the practice.

Using this information, answer the following questions.

- a. How much revenue will Monica be missing out on each year?
- b. How much after-tax operating profit will Monica be missing out on each year?
- c. How much free cash flow will Monica be missing out on each year?
- d. What is the cost of her hiatus in today’s dollars?
- e. If she decides to take a two-year hiatus starting next year, what is the cost at that point in time?
- f. If she decides to take a one-year hiatus starting two years from today, what is the cost at that point in time?
- g. She’s decided that her family cannot afford to miss out on more than \$3 million of value. What is the lowest opportunity cost of capital to ensure the cost of her hiatus does not surpass that threshold?

5.26 (*DCF, valuation*) Umbria Inc. is a healthcare insurer covering 3.4 million individuals in the U.S. Umbria is considering covering telehealth medical visits, in which patients engage with health care providers real-time via telephone and live audio-video with smartphone, tablet, or computer. The benefits of telehealth include:

- Reduced travel costs and inconvenience, increased access, and improved health care outcomes for patients.
- Reduced costs to the healthcare system.

The costs of telehealth include

- Increased technology costs, greater disruptions in continuity of care, and the lack of physical examination for patients.
- Overuse of medical services.
- Regulatory barriers.

The current estimated cost to Umbria of a telehealth visit is on average \$40 versus \$145 for in-person acute care - a substantial savings considering their policyholders currently visit their medical providers 1.9 times per year, on average.

However, implementing telehealth coverage is not without risk. Specifically, telehealth is only 83% as effective as in-person care, meaning 17% of the time an additional visit is required following a telehealth visit that would not be required following an in-person visit. Umbria's cost of capital for the telehealth coverage is 12%.

Using this information, answer the following questions.

- a. What is the current total annual cost savings of telehealth coverage if all of Umbria's customers utilize telehealth instead of in-person services (ignore the risk of less effectiveness)?
- b. If the cost per visit of both telehealth and in-person care grow at 3% per year, what is the present value of the cost savings, again ignoring an risk of lower efficacy and assuming that all of Umbria's customers switch to telehealth?
- c. Because of the lower efficacy of telehealth, how many total additional visits are to be expected assuming all of Umbria's customers switch to telehealth? What is the total cost associated with these additional visits?
- d. Assuming the efficacy of telehealth remains constant and the costs of in-person care grow at 3% per year, what is the present value of the additional costs of telehealth's extra visits?

5.27 (*Taxes benefits, Conceptual*) Derive the semi-elasticity of net present value with respect to the marginal tax rate, where the semi-elasticity is defined as

$$\frac{\partial NPV}{\partial \tau} \times \frac{1}{NPV}.$$

Using your result, answer the following questions.

- a. What is the interpretation of the semi-elasticity?

- b. What is the sign of the semi-elasticity?
- c. What is the economic implication of your answer to the previous question?

5.28 (*DCF, Decision criteria*) Traxor has developed a new plumbing wrench - the P-wrench - to work with PEX tubing - a flexible pipe that has become a popular alternative to more expensive copper pipe. Traxor has spent \$1 million on research and development over the previous two years designing and testing the product. They've concluded that the P-wrench has a three-year life until changes in pipe technology will render the P-wrench obsolete.

To manufacture the P-wrench, Traxor will build a new production facility that will take two years to build and require outlays of \$1 million today, \$2 million next year, and \$0.5 million two years from today. The plant has an expected life of ten years and estimated salvage value of \$0.75 million. Traxor plans on selling the plant for its \$1.25 million at the end of the project (five years from today). Additionally, Traxor will need to purchase \$1.5 million of equipment one year from today to manufacture the P-wrench. This equipment has a four-year usable life and no salvage value given its specialization.

To build enthusiasm for the product, Traxor will spend \$500,000 on marketing and sales two years from today. One year later - three years from today - they intend to go to market and forecast sales of 500,000 units. Sales are expected to contract by 200,000 per year as the market saturates and obsolescence approaches. The unit cost is \$13.80 and the unit sales price is \$46, both of which are expected to remain constant for the life of the project. Ongoing marketing, sales, and administrative expenses are anticipated to be \$4.8 million during the first year of sales, and stepping down \$1 million each year thereafter.

Traxor's net working capital for the project is estimated at 30% of sales, their marginal tax rate is 21%, and their cost of capital is 15%.

Using this information, answer the following questions.

- Construct a pro forma P&L statement for the project. What are the annual sales, operating expenses, taxes, and net operating profit after taxes? What are the tax implications of negative operating profits and on what assumption do they depend?
- Construct separate depreciation schedules for the plant and equipment. What are the annual depreciation amounts? What are the after-tax liquidation values?

- Construct a working capital schedule. What is the annual working capital investment for the project? How much working capital is recovered at the end of the project?
- Construct a free cash flow schedule? What are the annual free cash flows for the project?
- What is the value of the project today? What is the NPV of the project? Is this a viable project according to the NPV criterion?
- Is it “safe” to use the IRR criterion to assess the viability of this project? Why or why not? If safe, what is the IRR of this project and is it a viable project according to the IRR criterion?
- What is the payback period of the project?
- Based on your answers to the three previous questions, would you recommend moving forward with this project? Is there any other analysis you might perform?

Chapter 6

Project Selection

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

This chapter builds on the previous by examining how firms choose among different projects. Specifically, we

- identify the limitations of using the IRR rule for decision making,
- introduce a new decision metric, the profitability index, to help when selecting among projects when facing constraints, such as a budget or headcount limit.
- examine the implications of comparing projects with different lifetimes, and
- introduce customer lifetime value (CLV) and show that it is little more than discounted cash flow analysis applied to each customer.

This and the previous chapter provide the tools necessary to make the most corporate financial decisions. Two other corporate decisions - financial policy and acquisitions - are more specialized and presented in the third part of the book.

6.1 Choosing Among Mutually Exclusive Projects

Firms are often faced with selecting among a number of **mutually exclusive** opportunities or situations in which they can only select one. How do they decide among these options? Let's answer this question with a hypothetical business case. In 2014, the asset manager Vanguard decided to overhaul its information technology (IT) infrastructure. To do so, it solicited bids from several service providers, each of which provided two distinct bids. Assume we've already done the due diligence of estimating the free cash flows associated with each bid, all of which are summarized in Table 1. Our job is to assess these bids assuming Vanguard's cost of capital is 12% and select the best one.

Bid	Cash Flows
Cisco 1	\$100 million up-front cost to Vanguard for services rendered. \$60 million of cost savings in each of the following three years.
Cisco 2	Same cost savings as Cisco 1, \$60 million each year over the next three years. Now, the costs to Vanguard for services rendered are spread out over time. \$20 million of the cost is required up-front, and \$35 million of the cost is incurred in each of the following three years.
Juniper 1	\$100 million up-front cost to Vanguard for services rendered. \$90 million, \$70 million, and \$5 million of cost savings in each of the following three years, respectively.
Juniper 2	Juniper will pay Vanguard \$50 million up-front as an incentive and guarantee \$75 million in cost savings at the end of the second year. Vanguard will have to pay Juniper \$60 million at the end of the first year.
Huawei 1	\$20 million up-front cost to Vanguard for services rendered. \$20 million of cost savings in each of the following three years.
Huawei 2	Huawei will pay Vanguard \$50 million immediately as an incentive and guarantee \$50 million in cost savings each year for the next two years. The cost to Vanguard for services rendered is \$125 million to be paid at the end of the third year, one year after the last cost savings are experienced.

Table 1: Vanguard IT Bids

6.1.1 Comparing Cisco 1 and 2

Let's start by focusing on Cisco's two bids. The timeline for Cisco 1 is displayed in figure 6.1.

Years	0	1	2	3
Cash flows (\$mil)	-100	60	60	60

Figure 6.1: Cisco 1 Bid Timeline

The NPV and IRR for Cisco 1 are

$$NPV = -100 + \frac{60}{(1 + 0.12)} + \frac{60}{(1 + 0.12)^2} + \frac{60}{(1 + 0.12)^3} = \$44.11 \text{ million, and}$$

$$0 = -100 + \frac{60}{(1 + IRR)} + \frac{60}{(1 + IRR)^2} + \frac{60}{(1 + IRR)^3} \implies IRR = 36.31\%.$$

Because the NPV is positive, the bid is viable according to the NPV criterion. Because the IRR is greater than the cost of capital, the bid is also viable according to the IRR criterion. The IRR produces the same conclusion as NPV for this bid because the cash flows adhere to the sign rule - all cash outflows before cash inflows. In a nutshell, Cisco 1 looks like a promising bid.

Now consider the second Cisco bid, Cisco 2, whose timeline is in figure 6.2. For this bid, we've broken the cash flows up into savings and costs, the sum of which are the (net) cash flows of the bid - what Vanguard will ultimately pay or receive.

Years	0	1	2	3
Savings (\$mil)		60	60	60
Costs (\$mil)	-20	-35	-35	-35
Cash flows (\$mil)	-20	25	25	25

Figure 6.2: Cisco 2 Bid Timeline

The NPV and IRR for Cisco 2 are

$$NPV = -20 + \frac{25}{(1 + 0.12)} + \frac{25}{(1 + 0.12)^2} + \frac{25}{(1 + 0.12)^3} = \$40.05 \text{ million, and}$$

$$0 = -20 + \frac{25}{(1 + IRR)} + \frac{25}{(1 + IRR)^2} + \frac{25}{(1 + IRR)^3} \implies IRR = 111.85\%.$$

The NPV is positive and the IRR is greater than the cost of capital implying that Cisco 2 is also a viable bid. The agreement between the NPV and IRR criterion is to be expected because the cash flows of Cisco 2 adhere to the sign rule.

However, notice the contradiction that arises when *comparing* the two Cisco bids. Cisco 1 has a larger NPV, but smaller IRR than Cisco 2. How can Cisco 1 create more value

(higher NPV) but offer a lower return on investment (lower IRR)? More importantly, on which criterion should we rely? Is Cisco 1 or Cisco 2 better from a financial standpoint?

We mentioned in the previous chapter that NPV always gives the right answer. By that argument, Cisco 1 is preferred. But, that's an unsatisfying answer so let's build some intuition for why by first comparing the scale of each investment. Cisco 2 is a much smaller investment - \$20 million versus \$100 million. We earn 112% per year on \$20 million with Cisco 2, but 36.31% per year on \$100 million with Cisco 1. The difference in NPVs is telling us that the larger return on the smaller investment (Cisco 2) is generating less dollars than the smaller return on the larger investment (Cisco 1). But, we care about dollars and more is better. To hammer this point home, which would we prefer: a 100% return on a one dollar investment, or a 1% return on a \$1 million investment? The latter is clearly preferable because we get more money - \$10,000 versus \$1.

A closer look at Cisco 2 reveals another perspective on why the NPV of Cisco 1 is larger. While the projected cost savings are the same across the two bids, the payments to Cisco are quite different. Cisco 1 requires Vanguard pay \$100 million up front. Cisco 2 spreads the payments out over the life of the project: \$20 million up front and then \$35 million per year for the next three years. So, Cisco 2 contains a loan relative to Cisco 1. In Cisco 2, Cisco lends Vanguard \$80 million today in return for the three \$35 million repayments. The interest rate, r , on this loan is

$$80 = \frac{35}{(1+r)} + \frac{35}{(1+r)^2} + \frac{35}{(1+r)^3} \implies r = 14.93\%.$$

Note, $r=14.93\%$ is just the internal rate of return, which we found using Excel's *IRR* function.

The interest rate on the loan from Cisco is 14.93%, larger than Vanguard's 12% cost of capital. Vanguard would never want to borrow money from Cisco at a cost of 14.93% when they can raise money in the capital markets (banks, bond investors, shareholders) at a cost of 12%. This high interest rate loan embedded in the bid explains why Cisco 2's NPV is lower than Cisco 1's NPV. The high interest rate loan destroys value for Vanguard whose opportunity cost is 12%. NPV recognizes this value destruction, IRR does the exact opposite. The smaller upfront investment increases the return on investment for Cisco 2.

Had the payments in Cisco 2 been lower, say \$33.31 million per year instead of \$35 million, the loan interest rate would have been 12%, exactly equal to the cost of capital. Consequently, the NPVs of Cisco 1 and Cisco 2 would be identical. Had the loan interest rate been lower than 12%, the NPV for Cisco 2 would have been higher than that for Cisco 1. This example shows how finance can create or destroy value, and why the CFO's job of finding low cost financing is so important to the company.

Figure 6.3 illustrates the relation between the two Cisco bids. The horizontal axis measures the discount rate, r . The vertical axis measures NPV. The IRR for each bid is indicated by the circles on the horizontal axis - the discount rates that make the NPV of each bid equal to zero. Vanguard's cost of capital is indicated by the vertical, black dashed line at 12%.




Figure 6.3: NPV-Discount Rate Relation for Cisco 1 and 2 Bids

Both lines slope down reflecting the negative relation between the cost of capital and NPV - higher cost of capital, lower value because the project is riskier. The two lines intersect when the cost of capital is 14.93%, exactly equal to the loan rate in Cisco 2. At this point, Vanguard is indifferent between the two bids because they offer the same NPV (\$37.14). When Vanguard's cost of capital is below 14.93%, Vanguard prefers Cisco 1 because it can raise money at a lower cost than the rate at which Cisco is lending. The Cisco 1 bid provides more value as suggested by the blue curve being above the red curve. When Vanguard's cost of capital is above 14.93%, Vanguard prefers Cisco 2 because Cisco is lending at a lower interest rate than the rate at which Vanguard can raise money. It can raise money at a lower cost than the rate at which Cisco is lending. The Cisco 2 bid provides more value as suggested by the red curve being above the blue curve.

Figure 6.4 zooms in on figure 6.3 and adds two new lines - the dashed green and purple lines. By altering the loan rate implicit in their second bid, Cisco can alter the cash flows and the NPV of the bid. The green dashed line shows what happens when the loan rate is reduced to 6.13%. The Cisco 2 bid is now more valuable than Cisco 1 - the green dashed line is above the blue line for all cost of capitals above 6.13%. Intuitively, Cisco is lending money to Vanguard at a cost, 6.13%, well below their cost of capital, 12%. Similarly, if Cisco

increases the loan rate to 23.38% - the purple dashed line - the bid becomes less valuable because the loan rate is greater than Vanguard's cost of capital.




Figure 6.4: NPV-Discount Rate Relation for Cisco 1 and 2 Bids

To summarize, Cisco is providing financing to Vanguard with their second bid. The benefits are otherwise similar. So, the question becomes: How does the cost of financing compare to Vanguard's cost of capital? When the financing cost is greater than Vanguard's cost of capital, this financing destroys value. When the financing cost is less than Vanguard's cost of capital, this financing creates value.

Of course, this is another example of a theme we have seen repeatedly throughout the book. Vanguard has money and the question is what do they do with it? They can spend the money on bid 1 or they can borrow money from Cisco and spend the money elsewhere to earn their opportunity cost, 12%. Because Cisco is charging them 15%, more than their opportunity cost, it makes sense for them to not take that high interest rate loan. Choose Cisco 1.

This situation is conceptually no different than the decision to pay cash or borrow money for a car, the challenge Professor Keppel faced back in chapter 3. We saw that because Keppel's opportunity cost was below the loan interest rate, taking out the loan was a bad idea. It's the same thing here. Cisco is offering Vanguard a loan at 14.93% when Vanguard's opportunity cost is 12%. The loan is a bad idea. If the loan interest rate was reduced below 12%, it would then become a good idea for Vanguard.

We might be thinking: "Wait. Cisco 1 requires a \$100 million investment today. Cisco 2 only requires \$20 million, leaving \$80 million of dry powder that can be invested in other projects!" Put differently, there is an opportunity cost to spending an extra \$80 million on

Cisco 1. In fact, we know the opportunity cost of that \$80 million dollars. It's Vanguard's 12% cost of capital. If Vanguard wants to invest in other projects after investing in Cisco 1, they can simply raise more money at a cost of 12%.

Now we might be thinking that Vanguard can't raise more money because of budgeting constraints or that Vanguard's cost of capital will increase after spending so much money. This may be true but if so we've changed the rules of the game. By saying there is a limit to how much money Vanguard can spend or raise, and in reality this is true, we are introducing **financial constraints** or limitations on the firm's ability to raise capital. We'll get to this shortly. For now, as long as that the cost of capital is unaffected by the choice of project and Vanguard can continue to raise capital for other projects regardless of which bid they choose, Cisco 1 is unambiguously better than Cisco 2.

6.1.2 Juniper 1 and 2

Now let's compare Juniper 1 and 2, which are summarized in table 2. Juniper 1 is similar to Cisco 1 in terms of cost - \$100 million today. However, the cash flows are more front-loaded. The result is a higher IRR, 42%, relative to Cisco 1 but a lower NPV, \$39.72 million because of the relatively small inflow in year three.

Bids	Period				NPV	IRR
	0	1	2	3		
Juniper 1	-100.0	90.0	70.0	5.0	39.72	41.84%
Juniper 2	50.0	-60.0	75.0	0.0	56.22	#NUM!

Table 2: Juniper 1 and 2 Bids (\$mil)

Juniper 2 is more interesting. Vanguard will receive an upfront cash incentive of \$50 million in exchange for a \$60 million cost in year one and then savings of \$75 million in year two. The NPV of this bid is \$56.22 million, the highest thus far. However, when we compute the IRR Excel returns "#NUM!". Figure 6.5 illustrates what's happening. Juniper 1, represented by the green line, crosses the horizontal axis at 41.84% - the bid's IRR. However, Juniper 2 never crosses the horizontal axis meaning this bid doesn't have a "real" IRR.

Juniper 2 has two "imaginary" or "complex" IRRs, $(0.17 - 0.80i)$ and $(0.17 + 0.80i)$, where "i" equals the $\sqrt{-1}$. Complex numbers are way, way beyond the scope of this book and don't have a meaningful financial interpretation. Plus, good luck explaining them to

most investors or managers. This outcome is a result of two features of Juniper 2, the cash flow signs and the magnitude of the cash flows. The cash flows change signs twice, first from positive to negative in years zero and one and then negative to positive in years one and two. Multiple sign switches like this result in multiples IRRs, sometimes real sometimes, sometimes not real.




Figure 6.5: NPV-Discount Rate Relation for Juniper 1 and 2 Bids

6.1.3 Huawei 1 and 2

Table 3 details Huawei's cash flows and decision metrics. Huawei 1 has a similar upfront cost to Cisco 2, \$20 million, but the savings are less, \$20 million instead of \$25 million. Consequently, the NPV of Huawei 1, \$28.04 million, is less than that of Cisco 2. Its NPV also less than that of every other bid, though its IRR is over 80%.

Bids	Period				NPV	IRR
	0	1	2	3		
Huawei 1	-20.0	20.0	20.0	20.0	28.04	83.93%
Huawei 2	50.0	50.0	50.0	-125.0	45.53	-8.84%

Table 3: Huawei 1 and 2 Bids (\$mil)

Huawei 1 can also be viewed relative to Cisco 1, as illustrated in Figure 6.6. Huawei 1 generates the same cost savings as Cisco 1, \$60 million per year and only costs \$20 million upfront. In return for this \$80 million loan - remember Cisco 1 costs \$100 million upfront -

Vanguard must pay Huawei \$40 million per year leaving a net cash flow of \$20 million per year. The interest rate on this loan can be found as the IRR of the loan cash flows.

$$80 = \frac{40}{1+r} + \frac{40}{(1+r)^2} + \frac{40}{(1+r)^3} \implies r = 0.2338$$

The loan rate of 23.38% is almost double Vanguard's cost of capital, which is why the NPV of this bid is so low.

Years	0	1	2	3
Savings (\$mil)		60	60	60
Costs (\$mil)	-20	-40	-40	-40
Cash flows (\$mil)	-20	20	20	20

Figure 6.6: Huawei 1 Bid Timeline

Huawei 2 is a little less straightforward. The NPV is \$45.53 million, the second most valuable bid behind Juniper 2. However, the IRR is -8.84%. To help understand what's happening, Figure 6.7 illustrates the value of these two bids. Huawei 1 is represented by the downward sloping blue line that crosses the horizontal axis at 84% - its IRR. Huawei 2 is *upward* sloping and crosses the horizontal axis at -8.84% - its IRR.




Figure 6.7: NPV-Discount Rate Relation for Huawei 1 and 2 Bids

The upward slope implies that as the cost of capital increases, the value of the bid increases - seemingly the opposite of our intuition from the Fundamental Value Relation. However, if we look at the signs of the cash flows of Huawei 2, we see they are reversed in a sense. All the positive cash flows come before the negative cash flows. Huawei 2 is a

loan, albeit an unconventional one. Huawei is going to pay Vanguard \$50 million upfront and generate \$50 million in cost savings for Vanguard in return for a \$125 million repayment three years from today. When borrowing money, we like the lowest interest rate possible and a negative interest rate is as good as it gets. A negative interest rate implies that the lender is paying us to take their money. In this case, Vanguard is effectively being paid to take Huawei's money both directly through the upfront payment and through the cost savings their work will generate.

When assessing a project in which all the cash inflows come before all the cash outflows, the IRR decision rule is reversed. We should accept all projects in which the IRR is *less* than the cost of capital. Of course, this is just another way of saying that borrowing money is a good idea when the loan interest rate is less than our opportunity cost, what we discussed in detail in chapter 3.

6.1.4 Further Discussion

Capital Budgeting is Not a Zero-Sum Game

It's tempting to wonder why any company would bid on this project if all they're doing is creating lots of value for Vanguard at their own expense. Take Huawei 2, which is tantamount to a negative interest rate loan. Why would Huawei ever do such a thing? In other words, how could that possibly be profitable for Huawei? First, Huawei is not literally lending Vanguard money. We just recognized that the cash flows for the bid are analogous to a loan.

Second, these bids are not zero-sum games. They can be beneficial for both sides and, in fact, they must be beneficial for both sides otherwise one or both parties would never agree to the deal. Take Huawei 2. It must be the case that the cost of providing the service that generates the cost savings for Vanguard is less than the present value of the \$125 million they receive in three years. For example, it could cost Huawei \$80 million today at their cost of capital, say 10%, to install the new IT infrastructure. In this case, the NPV of the project to Huawei is

$$-80 + \frac{125}{(1 + 0.10)^3} = \$13.91 \text{ million.}$$

So, let's not confuse a gain for one side as a loss for the other when it comes to many business agreements.

Additional Shortcomings of the IRR and How to Address Them

We saw that Juniper 2 has no IRR. There are also situations in which there are multiple IRRs. Consider another bid by Huawei, Huawei 3, detailed in table 4. The IRR reported by Excel's IRR function is 157%. However, figure 6.8 shows that there are in fact three different IRRs: 157.5%, -32.5%, and -215.0%. (The figure is truncated along the vertical axis because the NPV gets so extremely large and small at the negative IRRs.) All three IRRs are equally valid and simply selecting the positive one, while convenient, is not justified on any economic grounds.

Bids	Period				NPV	IRR
	0	1	2	3		
Huawei 3	-10	21	20	-20	10.46	157.51%

Table 4: Huawei 3 Bid (\$mil)




Figure 6.8: NPV-Discount Rate Relation for Huawei 3 Bid

So, what do we do when there is no IRR or there are multiple IRRs. Fortunately, the cash flow streams that cause these problems - streams that violate the sign rule in which all outflows must precede all inflows - are far from the norm in practice. Additionally, this situation isn't really a problem because we should rely on NPV and not worry about the IRR. Nonetheless, some practitioners turn to a "modified" internal rate of return (MIRR), which is can be computed using the MIRR function in Excel. (See the technical appendix of this chapter for details and the Spreadsheet appendix for computation of MIRR in a spreadsheet.)

The MIRR will provide a single return estimate for any set of cash flows. However, to do so, it must modify the cash flows and make some assumptions. So, it is unclear exactly what an MIRR measures. Nonetheless, if a return on investment must be reported, the MIRR can come in handy, though its use as a decision making tool is unsupported by financial theory and is suspect. Use NPV.

Summary

This application demonstrated a number of shortcomings of IRR. Don't confuse these shortcomings as implying that IRR shouldn't be used in practice. It's too popular and unlikely to disappear anytime soon, if ever. Rather, this application should serve as a cautionary tale for using IRR as the *only* decision criterion. As suggested in the previous chapter, all three metrics - NPV, IRR, and payback period - should be presented in most business cases. However, NPV should be relied upon for decision making. Further, special care must be paid to situations in which IRR or payback period are nonsensical or conflict with NPV.

6.2 Choosing Among Projects when Facing Constraints

What happens when we have to make decisions when facing constraints or limitations? For example, we could have a financial budget limiting how much we can spend each period. Or, we could have a limited number of employees, warehouse space, or even time to perform tasks. Resources are finite, which limits our decisions.

Reconsider the Vanguard IT overhaul and imagine that our budget limits us to spending no more than \$20 million in any single year. In this case, the only projects that are feasible are Cisco 2 and Huawei 1, both of which require a \$20 million upfront investment. All the other projects have annual outflows exceeding our budget. Between these two projects, Cisco 2 has a higher NPV and should be chosen. Thus, the constraint, in this case a financial budget, leads to different decision than the unconstrained case.

Now let's consider a situation in which we have to choose among different projects that are no longer mutually exclusive - meaning we can choose more than one - and in which we face a budget or **financial constraint**. How can we choose the projects that will (1) generate the most value and (2) respect the budget (i.e., not require us to spend too much)?

6.2.1 Cypress Technologies

Cypress Technologies is a hypothetical company that creates digital devices, such as thermostats, pressure gauges, and monitoring sensors. The Chief Operating Officer (COO) is tasked with the development, production, and distribution for 6 different products whose upfront costs and NPVs are detailed in Table 5.

Product	Initial Cost	NPV
Thermo	128.0	360.0
Thermo Rx	74.0	59.2
Thermo Chem	25.0	90.0
SensorView	206.0	247.2
PressureView	350.0	1,365.0
PressureView+	97.0	485.0
Total	880.0	2,606.4

Table 5: Cypress Technologies Product Lineup (\$000s)

Without any constraint, the COO should undertake all six projects because each one is a positive NPV investment. However, the COO has a budget of only \$600,000. Thus, the question is: Which products should she undertake to maximize the value of her investments? Equivalently, what is the best allocation of her limited financial resources?

One option is to rank the projects by NPV and the highest ranked projects until the budget is exhausted. Table 6 presents the Cypress Technologies products ranked from high to low in terms of NPV. With a \$600,000 budget, the COO can undertake PressureView, PressureView+, and Thermo for a total NPV of $1.365 + 0.485 + 0.360 = \2.21 million.

Product	Initial Cost	NPV	Cum. Cost
PressureView	350.0	1,365.0	350.0
PressureView+	97.0	485.0	447.0
Thermo	128.0	360.0	575.0
SensorView	206.0	247.2	781.0
Thermo Chem	25.0	90.0	806.0
Thermo Rx	74.0	59.2	880.0
Total	880.0	2,606.4	

Table 6: Cypress Technologies Product Lineup Ranked by NPV (\$000s)

6.2.2 Profitability Index

One problem with simply selecting the highest NPV projects until we run out of money is that it ignores the constraining resource - money in this case - when ranking projects. To account for constraint, we can first scale our measure of value, NPV, by the constraining resource. Doing so leads to the **profitability index** (PI) defined as

$$\text{Profitability Index} = \frac{\text{Net Present Value}}{\text{Amount of Resource Consumed}}. \quad (6.1)$$

The profitability index measures the value per unit of resource or “bang for the buck” when the resource of interest is money.

Table 7 computes the PI for each product and then ranks products accordingly. The rightmost column presents the cumulative cost of the products moving down the table. The first product costs \$97,000, the first two products cost $97,000+350,000=\$447,000$, and so on. Ranking products by PI suggests that the COO should undertake four projects: PressureView+, PressureView, Thermo Chem, and Thermo.

Product	Initial Cost	NPV	Profitability Index	Cumulative Cost
PressureView+	97.0	485.0	5.0	97.0
PressureView	350.0	1,365.0	3.9	447.0
Thermo Chem	25.0	90.0	3.6	472.0
Thermo	128.0	360.0	2.8	600.0
SensorView	206.0	247.2	1.2	806.0
Thermo Rx	74.0	59.2	0.8	880.0

Table 7: Cypress Technologies Product Profitability Indices Against Budget (\$000s)

There are several interesting differences between the product rankings by NPV and PI that are detailed in Table 8. The product rankings are quite different. The most valuable project is PressureView with an NPV of \$1,365,000. The most valuable project *per dollar spend on the product* is PressureView+ with a PI of 5.0. Ranking products by NPV enables the COO to undertake three projects generating a total net value of \$2,210,000 and costing \$575,000. Ranking products by PI enables the COO to undertake four projects generating a total net value of \$2,300,000 and costing \$600,000. Ranking products by PI leads to a more efficient allocation of capital than ranking by NPV in this situation.

What happens if the constraining resource isn’t money? Imagine now that our COO has a large enough budget to undertake all six projects, but she only has 20 direct reports that

NPV Rank			PI Rank			
Product	NPV	Cum. Cost	Product	PI	Cum. Cost	NPV
PressureView	1,365.0	350.0	PressureView+	5.0	97.0	485.0
PressureView+	485.0	447.0	PressureView	3.9	447.0	1,365.0
Thermo	360.0	575.0	Thermo Chem	3.6	472.0	90.0
SensorView	247.2	781.0	Thermo	2.8	600.0	360.0
Thermo Chem	90.0	806.0	SensorView	1.2	806.0	247.2
Thermo Rx	59.2	880.0	Thermo Rx	0.8	880.0	59.2

Table 8: Cypress Technologies Product Profitability Indices Against Budget (\$000s)

can work on these projects? Table 9 provides the relevant information for this situation. The headcount column shows how many employees are needed for each project. The profitability index is NPV scaled by the required headcount. For example, SensorView generates a modest amount of value, \$247,200, but only requires 2 people. Consequently, it generates the most value per employee, \$123,600. The cumulative headcount column provides a running total of how many employees are required by the projects.

Product	Headcount	NPV	Profitability	Cumulative
			Index	Headcount
SensorView	2.0	247.2	123.6x	2.0
Thermo	3.0	360.0	120.0x	5.0
PressureView	12.0	1,365.0	113.8x	17.0
PressureView+	5.0	485.0	97.0x	22.0
Thermo Chem	1.0	90.0	90.0x	23.0
Thermo Rx	1.0	59.2	59.2x	24.0

Table 9: Cypress Technologies Product Profitability Indices Against Headcount (\$000s)

The decision making process here is the same as before. We've ranked projects by their profitability indices in Table 9 and selected the highest ranking projects until we exhaust our resource. In this case, the COO can undertake the SensorView, Thermo, and PressureView projects with her current headcount.

However, our COO is left with three extra employees that aren't assigned to a project. The COO has several options. She could assign two of the three to work on Thermo Chem and Thermo Rx, leaving her only one idle employee. Strictly speaking, doing so violates the profitability index rule by undertaking less valuable project; we're jumping over the PressureView+ project which has a higher PI. This example, therefore, highlights

a limitation of our profitability index rule, which is that the selection of projects must *completely exhaust the constraining resource* or else we may want to take on less valuable projects to fully utilize our resources. The COO she could fire the extra employees, but firing people, especially valuable employees, is costly. Finally, she could hire two additional people - consultants, part-time, full-time - so that she can undertake the PressureView+ project. The PI for the project, 97.0, tells us that she can compensate each additional employee an amount such that the present value of the after-tax compensation is \$97,000 plus any existing per employee cost. This is easier to see with an example.

The left side of Figure 6.9 presents the discounted cash flow analysis for the PressureView+ project as it was originally conceived: five employees working on the product that will generate \$485,000 of NPV, or \$97,000 per employee. (We've assumed there is no long-term investment or working capital requirements to avoid tangential details.) Decomposing the product NPV into costs and benefits, the present value of employee costs is \$34,710, or \$6,940 per employee. The present value of the employee benefits is \$519,710, or \$103,940 per employee. The difference between the employee costs and benefits is \$97,000 or the project PI.

The middle of figure 6.9 illustrates the maximum compensation to new employees before the project becomes value-destructive (i.e., negative NPV), assuming we keep the existing employees at their current compensation. We can pay each new hire \$249,440 in after-tax present value for the project to break even. While wildly lopsided, notice the average employee cost is \$103,940. Because the average employee generates \$103,940, the NPV per employee is \$97,000, our PI for the project.

The right side of figure 6.9 illustrates an alternative compensation policy that brings in two new hires and raises all existing employees compensation so that everyone is paid the same amount. Now the after-tax present value of the per employee costs are all equal to \$103,940, exactly what we found when we didn't increase the compensation of existing employees.

The message of this example is that the PI not only provides a useful decision tool when confronted with constraints, it also provides useful insights into the cost of relaxing the constraint. That is, how much we can afford to pay for more resources before the costs become so great that the project is no longer worth undertaking.

Summary

The **profitability index decision rule** is as follows.

Original Plan				New Hire Breakeven Compensation 1				New Hire Breakeven Compensation 2			
P&L	0	1	2	P&L	0	1	2	P&L	0	1	2
Sales		374.3	374.3	Sales		374.3	374.3	Sales		374.3	374.3
Emp 1		5.0	5.0	Emp 1		5.0	5.0	Emp 1		74.9	74.9
Emp 2		5.0	5.0	Emp 2		5.0	5.0	Emp 2		74.9	74.9
Emp 3		5.0	5.0	Emp 3		5.0	5.0	Emp 3		74.9	74.9
Emp 4		5.0	5.0	Emp 4		179.7	179.7	Emp 4		74.9	74.9
Emp 5		5.0	5.0	Emp 5		179.7	179.7	Emp 5		74.9	74.9
EBIT		349.3	349.3	EBIT		0.0	0.0	EBIT		0.0	0.0
Tax	0.2	69.9	69.9	Tax	0.2	0.0	0.0	Tax	0.2	0.0	0.0
NOPAT		279.5	279.5	NOPAT		0.0	0.0	NOPAT		0.0	0.0
Cost of capital	10%			Cost of capital	10%			Cost of capital	10%		
NPV	485.00			NPV	0.00			NPV	0.00		
PI(per emp)	97.00			PI(per emp)	0.00			PI(per emp)	0.00		
Employee costs	0	1	2	Employee 1-3 costs	0	1	2	Employee costs	0	1	2
Wages		25.0	25.0	Wages		15.0	15.0	Wages		374.3	374.3
Tax shield		5.0	5.0	Tax shield		3.0	3.0	Tax shield		74.9	74.9
After-tax wage		20.0	20.0	After-tax wage		12.0	12.0	After-tax wage		299.5	299.5
PV(Employee costs) per employee	34.71			PV(Employee costs) per employee	20.83			PV(Employee costs) per employee	519.71		
	6.94				6.94				103.94		
Employee benefits	0	1	2	Employee 4 & 5 costs	0	1	2				
Sales		374.3	374.3	Wages		359.3	359.3				
Tax		74.9	74.9	Tax shield		71.9	71.9				
After-tax sale		299.5	299.5	After-tax wage		287.5	287.5				
PV(Employee benefits) per employee	519.71			PV(Employee costs) per employee	498.88						
	103.94				249.44						
Total NPV	485.00			Total employee costs per employee	519.71						
NPV per employee	97.00				103.94						
				Employee benefits	0	1	2				
				Sales		374.3	374.3				
				Tax		74.9	74.9				
				After-tax sale		299.5	299.5				
				PV(Employee benefits)		519.71					
				Total NPV per employee		0.00					
						0.00					

Figure 6.9: Employee Compensation Scenarios for PressureView+ DCF

1. Compute the profitability index for each project by scaling NPV by the constraining resource (e.g., money, headcount, space).
2. Rank projects by their profitability index.
3. Select the highest ranking projects until the constraining resource is exhausted.

As we've seen, this recipe rests on a couple of assumptions. First, we have to completely exhaust the constraining resource as we move down the list, otherwise there may be projects with lower PIs that we want to undertake because they require fewer resources. Second, there is only one constraint. But, what if our COO faced both headcount and financial constraints? To solve precisely, more advanced mathematical techniques, such as integer and numerical programming, are required. There is software that makes solving these problems relatively simple, but this material is beyond the scope of this book and, unfortunately, not often used

by practitioners.¹ One practical, albeit imperfect, solution to this problem is to apply the profitability index decision rule using the most binding or important constraint.

6.3 Choosing Among Projects with Different Lifetimes

Roarke Stone Inc. uses a machine employing high pressure water jets to precisely cut stone for use in bathrooms and kitchens. The problem is that the high pressure water in conjunction with its continual use creates a great deal of wear and tear on the machine. Consequently, every few years Roarke is forced to replace the machine with a new one. It is currently choosing between two machines, which we will refer to as A and B.

- a Machine A lasts for four years and costs \$800,000. In the first three years of use, machine A will generate \$500,000 of free cash flow per year. In the fourth year, it will generate \$250,000 of free cash flow as a result of the machine wear.
- b Machine B lasts for three years and costs \$1,000,000. It will generate \$600,000 of free cash flow per year over its three-year useful life.

At the end of their useful lives, each option must be replaced with a new machine of the same type to avoid excessive retrofitting and training costs. Table 10 details the cash flows for each machine over one lifetime.

	Machine A	Machine B
Year	Cash Flows (\$)	Cash Flows (\$)
0	-800,000	-1,000,000
1	500,000	600,000
2	500,000	600,000
3	500,000	600,000
4	250,000	0

Table 10: Machine A and B Free Cash Flows Over One Life Cycle

Assuming Roarke's cost of capital for the machines is 10%, computing the NPV of the two machines is straightforward.

$$\begin{aligned} \text{Machine A NPV} &= -800,000 + \frac{500,000}{0.10} (1 - (1 + 0.10)^{-3}) + \frac{250,000}{(1 + 0.10)^4} = \$614,179.36 \\ \text{Machine B NPV} &= -1,000,000 + \frac{600,000}{0.10} (1 - (1 + 0.10)^{-3}) = \$492,111.19 \end{aligned}$$

¹Textbook references on these subjects include "Linear Optimization" by Dimitris Bertsimas and John N. Tsitsiklis and "Optimization Over Integers" by Dimitris Bertsimas.

Comparing these NPVs suggests that machine A is more valuable; it has a higher NPV. This conclusion would be true if we were only interested in purchasing a machine for one life cycle. However, because we will need to commit to a machine for an indefinite number of years, this comparison can be misleading. To compare the value of these two machines over multiple life cycles we need to standardize their time scales, that is, put them on an equal footing in terms of the lifetime over which we compare them. There are several different ways to do this.

6.3.1 Annuity Equivalent Cash Flow

The first approach is to compute an **annuity equivalent cash flow** that compares machine values over a single period. The annuity equivalent cash flow is the constant, periodic cash flow that generates the net present value for each machine. To compute it, we use our annuity cash flow formula, equation 2.2.

$$\begin{aligned} \text{Machine A annuity equivalent cash flow} &= \frac{614,179.36 \times 0.10}{(1 - (1 + 0.10)^{-4})} = \$193,755.66 \\ \text{Machine B annuity equivalent cash flow} &= \frac{492,111.19 \times 0.10}{(1 - (1 + 0.10)^{-3})} = \$197,885.20 \end{aligned}$$

These results imply that purchasing machine A is equivalent to receiving \$193,755.66 every year for four years starting next year, whereas purchasing machine B is equivalent to receiving \$197,885.20 every year for three years starting next year. Because the machines are replaced at the end of their life cycle, we will receive these cash flows for the indefinite future. We'd rather receive \$197,885.20 than \$193,755.66 each year so machine B is preferred to machine A.

Alternatively, we can find the least common multiple for the two lives, 12 years in this case, and compute the NPVs for each machine over this common lifetime. The cash flows over this horizon are detailed in table 11, which illustrates the problem of comparing the machines over one life-cycle. Machine A offers lower annual benefits, but Machine B comes with higher and more frequent acquisition costs. Which effect dominates can't be easily inferred from a comparison over just one life cycle. The NPVs of machine A and B over 12 years are \$1,320,191.33 and \$1,348,328.74, respectively. Again, machine B is preferred.

Of course, we could have arrived at the same estimates by discounting the annuity equiv-

	Machine A	Machine B
Year	Cash Flows (\$)	Cash Flows (\$)
0	-800,000	-1,000,000
1	500,000	600,000
2	500,000	600,000
3	500,000	-1,000,000+600,000
4	-800,000+250,000	600,000
5	500,000	600,000
6	500,000	-1,000,000+600,000
7	500,000	600,000
8	-800,000+250,000	600,000
9	500,000	-1,000,000+600,000
10	500,000	600,000
11	500,000	600,000
12	250,000	600,000

Table 11: Machine A and B Free Cash Flows Over Least Common Multiple Lifetime

alent cash flows over a 12-year period, as the following calculations show.

$$\text{Machine A NPV} = \frac{193,755.66}{0.10} (1 - (1 + 0.10)^{-12}) = \$1,320,191.33$$

$$\text{Machine B NPV} = \frac{197,885.20}{0.10} (1 - (1 + 0.10)^{-12}) = \$1,348,328.74$$

Thus, in practice, we only need to compute the annuity equivalent cash flow when comparing machines of different lives because this cash flow can then be used to compute the NPV over any horizon.

6.3.2 When to Invest

Now consider what would happen if Roarke had an existing machine in place that would serve for another three years and generate the following cash flows: \$500,000 one year from today, \$200,000 two years from today, and \$100,000 three years from today. Based on the analysis above, Roarke knows it wants to purchase machine B. The decision it faces now is *when* to purchase it: Today? One year from today? Two years from today?

We can quickly answer this question with what we've learned, but it's insightful to tackle the problem by systematically comparing the alternatives using first principles. Figure 6.10

presents the four alternatives under the Old machine's free cash flows. Each alternative corresponds to a different time at which to replace the existing machine. Replacing the existing machine today with machine B means Roarke will not receive any of the old machine's cash flows. Instead, it will face a cost of \$1 million today followed by \$0.6 million of cash inflows over the next three years until they replace machine B with a newer version and the cash flows repeat. Our analysis above showed that each life cycle of machine B is equivalent to receiving \$197,885.22 each year starting the year after replacement. Assuming the machines are replaced indefinitely, we can value this cash flow stream using our perpetuity formula (equation 2.3 to get $197,885.22/0.10 = \$1,978,852.0$). That is, replacing the machine today has a NPV of \$1.979 million.

Old machine	0	1	2	3	4
Free cash flows		500,000.0	200,000.0	100,000.0	0
Replace now	0	1	2	3	4
Existing FCF	0				
PV	0.0				
Annuity equivalent		197,885.2			
PV perpetuity	1,978,852.0				
NPV	1,978,852.0				
Replace 1-year	0	1	2	3	4
Existing FCF		500,000.0			
PV	454,545.5				
Annuity equivalent			197,885.2		
PV perpetuity	1,798,956.3				
NPV	2,253,501.8				
Replace 2-year	0	1	2	3	4
Existing FCF		500,000.0	200,000.0		
PV	619,834.7				
Annuity equivalent				197,885.2	
PV perpetuity	1,635,414.8				
NPV	2,255,249.6				
Replace 3-year	0	1	2	3	4
Existing FCF		500,000.0	200,000.0	100,000.0	
PV	694,966.2				
Annuity equivalent					197,885.2
PV perpetuity	1,486,740.8				
NPV	2,181,707.0				

Figure 6.10: Roarke's Existing Machine Replacement Alternatives

If Roarke replaces the machine one year from today, they receive \$500,000 from use of the existing machine, which has a present value of $500,000/(1+0.10) = \$454,545.5$. From year 2 onward, Roarke receives the annuity equivalent cash flow from machine B, \$197,885.2,

which in perpetuity has a present value of

$$\frac{197,885.2}{0.10} \times \frac{1}{(1 + 0.10)} = \$1,798,956.3$$

The perpetuity formula gives us the value as of year 1, so we need to discount that value back one year to get the value as of today. Adding the present value of the cash flow from the existing machine with the present value of the cash flows from the new machines gets us the value of replacing the machine one year from today, $454,545.5 + 1,798,956.3 = \$2,253,501.8$. We can already see that waiting one year to replace the existing machine is preferable to replacing it today.

Replacing the machine two years from today means we'll receive the next two cash flows generated from the old machine, \$500,000 and \$200,000. The sum of the present values of these cash flows is \$619,835. The present value of the annuity equivalent cash flow thereafter is

$$\frac{197,885}{0.10} \times \frac{1}{(1 + 0.10)^2} = \$1,635,415.$$

The NPV of replacing the machine two years from today is therefore $619,835 + 1,635,415 = \$2,255,250$, slightly larger than the NPV from replacing the machine one year from today.

Estimating the value of replacing the machine three year from today proceeds similarly. Looking across all four options reveals that the most valuable is replacing the machine in two years. However, there is a more efficient approach to arriving at this answer. We want to replace the existing machine before the annuity equivalent cash flow of the new machine is greater than the cash flows generated by the existing machine. In other words, keep using the existing equipment as long as it is generated cash flows larger than what can be generated by the new machine on an annuity equivalent basis. The annuity equivalent cash flow is 197,885.2, which is less than the 500,000 and 200,000 the existing machine generates over the first two years. Only in the third year does the cash flow from the existing machine, 100,000, fall below that threshold implying we should replace the existing machine before experiencing that lower cash flow.

Of course, this rule only works if the cash flows from the existing machine are declining over time, which is often the case with fixed capital, and the costs and benefits of the replacement machine is not changing as we delay investment. Were either of these scenarios true, our first principles approach, with appropriate cash flows, illustrated in Figure 6.10 will always work.

6.4 Customer Lifetime Value (CLV)

In the late 1980s, marketing executives realized that they could apply discounted cash flow analysis to customers. Effectively, each customer is treated as an individual project generating a sequence of cash flows that can be discounted and summed to estimate a **customer lifetime value (CLV)** or just **lifetime value (LTV)**. Further, by summing the CLVs of all our customers, we could in theory get the value of a project, division, or even an entire company.

While there is debate over what aggregating individual customer values actually measures, the explosion of customer level transaction data and subscription-based business models has made customer lifetime value a key financial metric for many businesses and investors. Let's illustrate its implementation and introduce some of the associated lingo; marketers just had to have their own jargon. CLV analysis comes in a variety of different versions varying in complexity. We'll start with the basics and add a few bells and whistles as we go. The goal isn't to cover the entirety of CLV analysis, only to show that it is nothing more than DCF in disguise and often makes some assumptions that are financially questionable.

6.4.1 The Basic Model

Perhaps the most basic version of a customer lifetime model is the following.

$$CLV = \text{Duration} \times \text{Recurring revenue} \times \text{Operating margin} \quad (6.2)$$

Duration is how long the customer sticks around and spends money. The recurring revenue is how much they spend each period (e.g., month, year). Operating margin is the fraction of the revenue the company keeps after netting out expenses.

For example, consider a SaaS (Software as a Service) hypothetical company we'll call Cuesta focused on business to business (B2B) engagements. Cuesta offers business customers financial planning and analysis (FP&A) software in the cloud. For a recurring fee, customers can access this software remotely to perform their financial analytics. The typical customer engages with Cuesta for approximately three years by paying a monthly fee of \$600. This fee is revenue for Cuesta and is referred to as a **monthly recurring revenue** or **MRR**.

Cuesta's direct costs and overhead for each customer are \$120 per month implying an operating margin of 80%. These costs include customer support, product development, and fees Cuesta pays to Amazon Web Services (AWS), which hosts Cuesta's software and its

clients' data on their cloud platform. The CLV for a Cuesta customer in this scenario is

$$CLV = 36 \text{ months} \times \frac{\$600}{\text{month}} \times 0.80 = \$17,280.$$

Often analysts will factor in the money spent on acquiring the customer through an additional term referred to as **customer acquisition costs** or **CAC**. These costs include marketing, advertising, sales expenses, etc.

$$CLV = -CAC + \text{Duration} \times \text{Recurring revenue} \times \text{Operating margin} \quad (6.3)$$

For our example, let's assume the CAC is \$5,000 per customer, implying a customer lifetime value net of customer acquisition costs of $17,280 - 5,000 = \$12,280$.²

Some common key performance indicators for CLV analysis include.

1. LTV-to-CAC Ratio

$$\text{LTV-to-CAC} = \frac{17,280}{5,000} = 3.46x$$

Each customer generates \$3.46 for each dollar we spend to acquire them.

2. CAC Payback Period

$$\text{Payback Period} = \frac{CAC}{\text{Recurring revenue} \times \text{Operating margin}} = \frac{5,000}{600 \times 0.8} = 10.4 \text{ months}$$

It will take 10.4 months in to recover the money spent acquiring the customer (CAC) from the profits we receive from them.

3. Churn Rate

$$\text{Churn} = \frac{1}{\text{Expected customer lifetime}} = \frac{1}{36} = 0.0278$$

Every month approximately 2.8% of the customer base will leave the company, *assuming* signups and exits are uniformly distributed throughout the year and all customers behave similarly with regard to when they leave Cuesta. The churn rate tells us how many customers the company needs to sign up just to keep the customer base - and therefore revenue stream - constant.

²Whether the CAC is treated as an operating expense or capital expenditure depends on the nature of the investment and accounting rules. This distinction could be important for tax reasons.

6.4.2 DCF in Disguise

Though it may sound different, customer lifetime value and most of its associated KPIs are just familiar financial concepts. The LTV-to-CAC ratio is just a multiple on invested capital or MOIC. The CAC payback period is our familiar payback period. And, CLV is discounted cash flow analysis. To see this last point, let's write the CLV expression (equation 6.3) in a slightly different manner.

$$\begin{aligned}
 CLV &= \underbrace{\text{Customer Acquisition Costs}_0}_{CashFlow_0} \\
 &+ \underbrace{(\text{Sales}_1 - \text{Operating expenses}_1)}_{CashFlow_1} + \underbrace{(\text{Sales}_2 - \text{Operating expenses}_2)}_{CashFlow_2} + \dots \\
 &+ \underbrace{(\text{Sales}_T - \text{Operating expenses}_T)}_{CashFlow_T}
 \end{aligned}$$

When viewed like this, we can recognize equation (6.3) as net present value in which we have assumed:

1. Sales and expenses are constant;
2. No taxes;
3. No long-term investment;
4. No working capital investment; and,
5. The discount rate, r , is zero.

So, the accuracy of CLV depends on whether these assumptions make sense.

In our Cuesta example, it's possible that revenue and operating expenses are constant because of the nature of the service. Though, varying usage of services might generate varying costs. It's also possible that long-term investment is zero because Cuesta is outsourcing its IT needs to Amazon and may rent office space. Taxes could also be zero if Cuesta is unprofitable or domiciled in a tax-free jurisdiction.

However, no investment in net working capital is a stretch unless customers pay cash every month. This seems unlikely in our Cuesta example or even a business to customer (B2C) setting in which customers pay by credit card. Upfront payments for extended service periods, e.g., one year, also generate current liabilities that must be recognized. Perhaps most troubling is the lack of discounting. Even if the cash flows are near risk-free because the

relationship is governed by a contractual commitment, the discount rate shouldn't be zero. It should be the risk-free rate! (Yep, I'm getting worked up here...)

Of course, we don't need to make any of these assumptions in practice. With technology as of 2022, we can maintain customer level forecasts for each component of free cash flow: revenue, expenses, long-term investment, and working capital. We can also estimate customer level discount rates or, more realistically, market segment discount rates. So, don't view what is often employed in practice as anything but simplifying assumptions that can easily be relaxed if necessary.

The message here is that CLV is just DCF by a different name. Recognizing the similarities (and differences) is important because doing so makes explicit the assumptions behind any CLV analysis.

6.5 Key Ideas

Corporate financial decision making is fundamentally no different from personal decision making. We want to undertake decisions with positive net present value, decisions in which the present value of the benefits are larger than the present value of the costs. In a corporate setting, the cash flows used to measure the costs and benefits are called free cash flows and measured in a specific way (equation 5.3).

- Whether we are looking at an individual project or many projects, the NPV criterion will unambiguously identify those that create ($NPV > 0$) or destroy ($NPV < 0$) value.
- When assessing the viability of an *individual* project, the IRR criterion will produce the same outcome - accept or reject - as the NPV criterion as long as the cash flows adhere to the sign rule: all cash outflows (negative cash flows) occur before all cash inflows (positive cash flows).
- When selecting among projects, the IRR rule should *not* be relied upon even if the projects adhere to the sign rule.
- The profitability index (PI) can be used to select among project when facing constraints by dividing project NPVs by the constraining resource (budget, headcount, office space, materials, etc.). We can then rank projects by their PIs and the select projects from high to low PI until we exhaust the constraining resource.

- Two caveats to the PI recipe apply. First, we have to completely exhaust the resource, or else lower PI projects with lower resource requirements may be preferable. Second, the PI can only handle one constraint. Though, a practical, if imperfect, solution can be found by focusing on the most binding constraint.
- One-off projects with different lives don't create any complications for our NPV criterion - choose the project with the higher NPV. However, when assessing projects with different lives that are repeated, such as machine replacement, we have to be more careful. We discussed two approaches: compare the NPV over the least common multiple lifetime and compare the annuity equivalent cash flow.
- The timing of investments should be such that we continue to use existing machinery or equipment as long as they generate cash flows greater than the annuity equivalent of the new investment.
- Customer lifetime value (CLV) models are nothing more than discounted cash flow analysis with new labels. Be careful how you use them because they often hide important assumptions about both cash flow and discount rates.

6.6 Technical Appendix

The MIRR can be computed several different ways. The approach described here is the one implemented in Excel and is best illustrated with an example. Let's use the Juniper 2 bid from the Vanguard IT overhaul application. Table 12 details the cash flows and decision metrics assuming Vanguard's cost of capital is 12%.

Bids	Period				NPV	IRR
	0	1	2	3		
Juniper 2	50	-60	75		56.22	#NUM!

Table 12: Juniper 2 Bid

As long as there is only one sign change in the cash flow stream, we can always find a unique, real IRR. Juniper 2 has two sign changes: one moving from today (positive 50) to one year from today (negative 60), and one moving from one year from today (negative 60) to two years from today (positive 75).

The MIRR avoids this problem by first discounting negative cash flows (i.e., outflows) back to today at the cost of capital and compounding positive cash flows (i.e., inflows) to

the end of the project by a possibly different rate sometimes called the reinvestment rate. For this example, let's assume the reinvestment rate is the same as the opportunity cost of capital. The present value of the outflows and future value of the inflows for Juniper 2 are as follows.

$$\begin{aligned} \text{Present value of outflows:} & \quad \frac{-60}{1 + 0.12} = -53.57 \\ \text{Future value of inflows:} & \quad 50(1 + 0.12)^2 + 75 = 137.72 \end{aligned}$$

Next, we take the perspective that Juniper 2 only consists of two cash flows: one outflow today of \$53.57 and one inflow two years from today of \$137.72. The annual return associated with these cash flows is

$$MIRR = \left(\frac{\text{Future value of inflows}}{-\text{Present value of outflows}} \right)^{1/T} - 1 = \left(\frac{137.72}{53.57} \right)^{1/2} - 1 = 0.6034,$$

or 60.34%.

6.7 Problems

For all problems requiring calculation, it is strongly recommended - in many cases required - that a spreadsheet or other computing program be used.

6.1 (*IRR sign rule*) The following table presents cash flows for six projects.

Project	Year				
	0	1	2	3	4
A	-100	120	0	0	0
B	-15	12	-14	35	0
C	-50	10	20	30	40
D	-110	-40	-30	80	400
E	-100	0	0	0	247
F	500	180	300	-690	0

What are the IRR and NPV of each project if the cost of capital is 12%? For which projects can we be sure that the NPV and IRR decision rules will lead to the same conclusion?

6.2 (*IRR rule intuition*) The follow table presents cash flows for three projects.

Project	Year			
	0	1	2	3
A	-100	50	40	30
B	-100	100	100	-80
C	100	50	60	-70

What are the IRR and NPV of each project if the cost of capital is 10%? Create a plot with discount rate on the horizontal axis and NPV on the vertical axis. Plot the discount rate-NPV relation for each project by varying the discount rate from -50% to 450%. What does the plot reveal?

- 6.3 (*NPV, IRR, and financing*) Exxon-Mobil is considering the acquisition of a new oil field off the coast of Brazil from the Brazilian government. The estimated cash flows of the deal are presented in the table.

	Year			
	0	1	2	3
Cash flows (\$mil)	-250	125	125	125

Exxon-Mobil has also been given the option to finance the deal through the Brazilian central bank. Specifically, Exxon-Mobil would pay \$150 million today, instead of \$250 million, and then pay \$26 million each year thereafter for *five* years. The future benefits of the acquisition, \$125 million for three years, would remain unaffected by the financing. (All cash flows are in U.S. dollars.) The project cost of capital is 12%.

Using this information, answer the following questions.

- What are the IRR and NPV of the project if Exxon-Mobil chooses *not* to finance the deal through the Brazilian central bank?
- What are the IRR and NPV of the project if Exxon-Mobil chooses to finance the deal through the Brazilian central bank?
- What is the implied interest rate on the loan in the financing package? What is the implied credit spread on the loan if the current yield on a five-year Treasury note is 6.2%?
- What is the NPV of the loan in the financing package? How is this value related to the NPV of the project with and without financing? To what does the NPV of the loan correspond?
- Should Exxon-Mobil accept Brazil's financing package? Explain why or why not.

- f. What annual loan repayment amount would make Exxon-Mobil indifferent between accepting and rejecting the financing package? What is the corresponding implied interest rate?

6.4 (*NPV, IRR, profitability index*) Hirschfield Enterprises is a chemical manufacturing company considering the production and distribution of two new compounds whose costs and benefits in \$millions are detailed in the following table.

Project	Year					
	0	1	2	3	4	5
Exotherm	(50.00)	45.00	35.00	25.00	15.00	0.00
Caprex	(250.00)	175.00	135.00	95.00	55.00	15.00

The project cost of capital for both projects is 8%, and only one of the projects may be selected.

Using this information, answer the following questions.

- What are the IRR, NPV, and profitability index (PI) for each project?
 - Which project is preferable according to the IRR criterion? NPV criterion? PI criterion? Which project should Hirschfield choose and why? Do your findings highlight any similarities between the IRR and PI criterion?
- 6.5 (*NPV, IRR, profitability index, financial constraints*) Rob Low, the CEO of SubX Maritime Inc., is considering how to spend his \$120 million annual budget. He is considering deploying four boats whose costs and benefits in \$millions are detailed in the following table.

Boat	Year			
	0	1	2	3
Delaware	(120.00)	90.00	80.00	70.00
Montana	(40.00)	30.00	25.00	20.00
Oregon	(50.00)	20.00	35.00	50.00
Vermont	(30.00)	10.00	25.00	40.00

The cost of capital for each boat is 15%.

Using this information, answer the following questions.

- What are the IRR, NPV, and profitability index (PI) for each project?

- b. Which boats should Rob deploy? Explain your reasoning. How much value will his choices generate for the company?
- c. If Rob's budget was \$250 million, which boats should he deploy?
- d. If Rob's budget was \$100 million, which boats should he deploy? What should he do with any excess funds?

6.6 (*DCF, decision criteria, Lifetime value (LTV), customer lifetime value (CLV)*) You are evaluating a Software as a Service (SaaS) business model put forward by a startup company, Cortend Partners. You are given the following information: Customer acquisition costs (CAC) of \$500 per customer are incurred one month prior to customer sign-up. You may assume that customers sign-up in period 1. These costs are comprised of sales and marketing expenses.

Each customer pays a monthly subscription fee of \$100 with the first payment due immediately upon sign-up. Customers can quit the service at any point in time at no charge and each customer has an expected lifetime of 12 months. (They stay on the platform and pay fees for 12 months.)

Monthly support costs equal 40% of monthly recurring revenue and are fully expensed. The annual cost of capital is equal to 20%. The business pays no taxes because of historical operating losses. Working capital requirements are negligible and can be ignored. The firm avoids any capital investment by having its software hosted in the cloud and expensing these costs, which are part of the recurring support costs.

Using this information, answer the following questions.

- a. Estimate the net operating profit after taxes (NOPAT) for the typical customer.
- b. Estimate the free cash flows for the typical customer.
- c. What is the typical customer's lifetime value (CLV), ignoring the time value of money and risk?
- d. What is the typical customer's lifetime value (CLV), accounting for the time value of money and risk? How does it compare to your answer in the previous problem?
- e. What is the net present value of a typical customer's cash flow stream?
- f. Estimate the lifetime value-to-CAC ratio using undiscounted cash flows? If the criterion for project acceptance is an LTV-to-CAC ratio greater than three, does this project get accepted or rejected?
- g. What is the internal rate of return on an average customer in APR terms, i.e., simple interest?

- h. To increase growth, the CFO suggests offering a subscription discount. At what monthly subscription fee will Cortend break-even on its customers?
- i. What is the lowest monthly recurring revenue Cortend can generate and still ensure an internal rate of return of at least 40% in APR terms, i.e., simple interest
- j. What is the internal rate of return (expressed as an APR) for an alternative customer segment who is otherwise identical except that its per customer acquisition costs are \$100 instead of \$500 and its monthly subscription fees are \$35 per month instead of \$100?
- k. True or False: Assuming the customer segments are mutually exclusive and of equal size, and you face no financing constraints, you prefer to target the customer segment described in the previous question.

6.7 (*DCF, decision criteria, financing*) HP Inc is offering to overhaul your Olivander's Wand Co's logistics with its new Wand technology featuring a Phoenix core. The cost of this overhaul is \$240 million today, with an annual increase in free cash flow over the next three years equal to \$100 million. Our company's cost of capital is 10% per annum.

Using this information, answer the following questions.

- a. Compute the NPV and the IRR of this logistics overhaul. Should you proceed with the overhaul? Explain your answer.
- b. Your company is working closely with the Private Equity Firm Snape Investments. According to Snape, you should ask for a payment plan with HP Inc. such that you do not pay the \$240 million up front but instead make 4 payments of \$70 million. The first payment happens today and the other three payments happen in years 1, 2 and 3. Snape argues that this will save a significant amount of money today, leaving your company with a lot of "dry powder" for other investments.
 - i. Compute the NPV and IRR of of the overhaul under this alternative financing scheme.
 - ii. What is the implied interest rate on the loan in this financing scheme?
 - iii. Do you think Snape's recommendation is financially sound? Explain.
 - iv. What size payments would leave you indifferent between paying the \$240 million up front and the equal payment financing strategy?

6.8 (*DCF, decision criteria, machine replacement*) Conrad Corp. is trying to determine the optimal replacement policy for one of its machines. The machine costs \$15,000,

has a usable life of three years, and a salvage value of \$3,000. The annual maintenance costs and corresponding liquidation values are detailed in the following table.

Year	Maintenance Costs (\$)	Salvage/Liquidation Values (\$)
1	1,000	6,000
2	2,000	3,000
3	3,000	0

The firm faces a 34% tax rate on all profits, a 12% project cost of capital, and uses a straight-line depreciation. The company's revenues are unaffected by the replacement policy and the company is profitable.

Using this information, answer the following questions.

- Compute the net present values for the three replacement times.
- Which replacement time offers the largest NPV? Is this the optimal replacement time?
- Estimate the per-year operating cost for each replacement time policy? Which offers the lowest cost? Does it align with your answer to the previous question?

6.9 (*DCF, decision criteria, machine replacement*) DMet Industrial is a metal shaper, manufacturing metal parts for use in the automotive sector. A key piece of equipment in the production of those parts is a lathe, which is used primarily for shaping the metal by rotating it around an axis. The lathe requires progressively more maintenance as it ages because of use. DMet has a schedule of annual maintenance costs and potential resale values for its existing lathe. These figures are detailed in the following table.

Year	Maintenance Costs (\$)	Resale Value (\$)
0	0	2,500
1	800	2,400
2	1,000	2,300
3	1,200	2,200
4	1,400	2,100
5	1,600	2,000

A new lathe will cost \$7,400 and require \$285 of annual maintenance that is realized beginning one year after purchase. The new lathe has a salvage and liquidation value

of \$2,800 at the end of its five-year useful life. DMet's cost of capital is 15%, and they face no taxes because of significant historical operating losses.

Using this information, answer the following questions.

- a. What is the optimal time to replace the machine?
- b. How does the optimal time to replace the machine change as the discount rate changes? Can you provide any intuition for your answer? (*Hint: You might consider constructing a data table in Excel*)
- c. How does the optimal time to replace the machine change as the new lathe maintenance costs change? Can you provide any intuition for your answer? (*Hint: You might consider constructing a data table in Excel*)
- d. How does the optimal time to replace the machine change as the new lathe salvage value changes? Can you provide any intuition for your answer? (*Hint: You might consider constructing a data table in Excel*)

- 6.10 (*DCF, decision criteria, investment timing*) Microsoft is considering a new project - Game Pass - by which gamers, via a subscription, can access hundreds of games in the cloud, as opposed to via DVDs sold through retailers. One implication of the Game Pass project is increased usage of Microsoft servers.

The increased usage will result in the acceleration of a plan to construct a new data center. Specifically, if the Game Pass project goes forward, the new data center will need to be built two years from today as opposed to the originally planned four years. The new data center is expected to cost \$850 million to construct and bring online. It will last for 10-years and cost \$40 million per year to operate beginning one year after construction. Microsoft's cost of capital is 8%.

What is the effect of this acceleration on the NPV of the Game Pass project? How would your answer change if every 10 years a new data center had to be constructed to deal with increasing demand assuming the construction and operating costs do not change.

- 6.11 (*Decision criteria, investment timing*) Trexor is a mining company that is deciding when to begin mining copper just outside of Sonora Mexico. While mining costs are likely to rise in the future, Trexor believes that dwindling supplies of copper and increasing demand will lead to increases in the price of copper that will more than offset these increasing costs.

To determine when to begin mining, Trexor computed the net present value of the mining operations for four different starting times. The NPVs, presented in the table, are *as of the date that mining operations begin*. If Trexor's cost of capital is 14% per year, when should they begin mining operations?

	Start Date of Mining Operations			
	0	1	2	3
NPV at start date of operations (time t)	230	272	301	339

6.12 (*Project selection, different length projects*) As part of a home renovation project, Michael is deciding what light bulbs to use in the recessed light fixtures. The table presents the options from which he is selecting.

Bulb	Type	Cost per bulb (\$)	Life (Years)	Annual Operating Costs (\$)
Soraa	LED	30.99	22	0.89
Sylvania	Halogen	8.63	4	4.93
Philips	Incandescent	2.08	2	5.91

Michael's annual opportunity cost of capital is 3%.

Using this information answer the following questions.

- What is the present value of the total costs - purchase plus operating - of one lifetime for each bulb?
- Using your answer from the previous question, what are the periodic equivalent cash flows for each bulb?
- Assuming Michael replaces the bulbs indefinitely, which bulb should he use? What is the present value of his cost savings relative to the other two bulbs if he needs 60 bulbs, again assuming the bulbs are used indefinitely (i.e., forever)?
- If Michael plans on flipping the house (i.e., selling) one year from now, which bulb should he choose? What are his total cost savings accounting for all 60 bulbs?
- (Advanced) What is the least amount of time Michael needs to stay in the house for it to be financially wise to use the LED bulbs relative to the Halogen bulb? Incandescent bulb? Create a line plot showing the relation between the time in the home in years and the total costs of operating each bulb. When is it optimal to use the Halogen bulb?

6.13 (*Project selection, different project durations*) This problem revisits the previous problem in a more precise manner. As part of a home renovation project, Michael is deciding what light bulbs to use in the recessed light fixtures. The table presents the options from which he is selecting.

Bulb	Type	Energy (Watts)	Cost per bulb (\$)	Life (Hours)
Soraa	LED	9	30.99	25,000
Sylvania	Halogen	50	8.63	4,000
Philips	Incandescent	60	2.08	2,000

Each bulb produces an equivalent amount of light despite different energy consumption. The local electricity rate is \$0.09 per kilowatt-hour implying that a bulb requiring 60 watts of energy would cost $0.09 \div 1000 \times 60 = \0.0054 to run for one hour. Daily usage for each bulb is estimated to be three hours.

Using this information answer the following questions.

- What are the hourly and daily operating costs for each light bulb?
- How many days and years will each bulb last?
- What is the present value of the total costs of one lifetime for each lightbulb? (Hint: Assume a period is one day.)
- Using your answer from the previous question, what are the periodic equivalent cash flows for each bulb?
- Assuming Michael replaces the bulbs indefinitely, which bulb should he use? What is the present value of his cost savings relative to the other two bulbs if he needs 60 bulbs, again assuming the bulbs are used indefinitely (i.e., forever)?
- If Michael plans on flipping the house (i.e., selling) one year from now, which bulb should he choose? What are his total cost savings accounting for all 60 bulbs?
- (Advanced) What is the least amount of time Michael needs to stay in the house for it to be financially wise to use the LED bulbs relative to the Halogen bulb? Incandescent bulb? Create line plot showing the relation between the time in the home in years and the total costs of operating each bulb. When is it optimal to use the Halogen bulb?
- How much additional insight is gleaned from the more precise calculations in this version of the problem?

6.14 (*Project selection, different project durations*) The Maybrook apartments has several vacant apartments that it must prepare and list for rent. Before renting a unit, the apartment must be prepped, a process that takes several days and requires a crew of maintenance workers. The preparation costs and crew sizes are detailed in the table, as are monthly rents. Maybrook management has found that different apartments cater to different segments of the population. These segments have different risk profiles in terms of their sensitivity to market conditions and expected tenancy. As such, listed with each apartment is a unique opportunity cost of capital and expected duration in the apartment.

Apartment	Prep Cost (\$)	Prep Crew Size	Monthly Rent (\$)	Discount Rate (%)	Tenancy Duration (Months)
1	5,000	3	2,800	8.0	18
2	4,500	3	2,795	9.0	18
3	1,000	2	1,400	15.0	12
4	10,000	4	4,700	3.0	24
5	3,000	2	1,600	12.0	12

Using this information, answer the following questions.

- Estimate the NPV, IRR, and multiple on invested capital (MOIC) for one tenant cycle for each apartment. (MOIC is defined as the total cash inflows divided by the total cash outflows ignoring any discounting.) Based on just one tenancy cycle, how would you rank the projects according to each metric? I.e., Which is the best project according to NPV? IRR? MOIC? Which is the worst? Discuss any differences in the rankings.
- Derive closed-form expressions (i.e., equation) for the payback period and the discounted payback periods and use your results to estimate these measures.
- Compute the profitability index of one tenant cycle for each apartment using prep cost as the denominator. Assuming management only has \$15,500 of cash on hand for prep costs, which apartments should it ready today according to your profitability index calculations?
- Compute the profitability index of one tenant cycle for each apartment using crew size as the denominator. Rank the apartments based on this measure. Assuming management only has ten maintenance workers available to prep apartments, which apartments should it ready today according to your profitability index calculations?

- (e) Rank the units in terms of their value to management assuming each apartment will be rented in perpetuity.

6.15 (*Capital budgeting, share repurchases*) As of December 2022, Alphabet Inc., the parent company of Google, had \$113.8 billion of cash and cash equivalents (e.g., money market funds, Treasury bills, commercial paper). In deciding what to do with that money, Alphabet management was considering opening a new data center, the details of which are as follows.

- The data center requires capital expenditures of \$100 million today, \$200 million next year, and \$50 million the year after in order to begin operating the center three years from today. The center's usable life is 10 years over which the capital expenditures will be straight line depreciated to a salvage value of \$35 million. The capital expenditures can only be depreciated *after* construction has been completed.
- Six months ago, Alphabet spent \$15 million determining the location for the data center.
- In its first year of operation - four years from today - the center is expected to generate \$100 million of revenue, which will grow by 10% per year over the center's life.
- Operating expenses - excluding depreciation - are a constant 60% of revenue. Additionally, Alphabet will spend \$10 million the year prior to the center opening to recruit employees.
- Net working capital, consisting of mostly cash and payables, is 22% of revenue generated by the data center, and is expected to be recovered in full shortly after the data center is sold.
- At the end of the data center's life, Alphabet anticipates selling the center for \$50 million, due largely to expected increases in land value.
- Alphabet's marginal tax rate is 18% and its cost of capital is 12% per annum.

Using this information, answer the following questions relating to the data center project.

- a. What are the revenue forecasts?
- b. What are the EBITDA forecasts?
- c. What are the depreciation forecasts?

- d. What are the EBIT forecasts?
- e. What are the NOPAT forecasts?
- f. What are the anticipated after-tax proceeds from the sale of the data center at the time of the sale (i.e., don't discount)?
- g. What is the annual net working capital required by the project?
- h. What are the free cash flow forecasts?
- i. What is the present value?
- j. What is the net present value?
- k. What is the internal rate of return, and does it make sense to use the internal rate of return criterion to assess the viability of the data center?
- l. What is the payback period?
- m. Alphabet is also considering repurchasing 10 million share at the current market price of \$88.73 per share. Assuming the shares are undervalued by \$5.00 per share, how much value can Alphabet create through this share buyback?
- n. Should Alphabet invest in the data center? Repurchase its shares? Both? Or, Neither?

6.16 (*Mortgage shopping*) Luna is preparing to buy a new home that costs \$850,000. She has \$250,000 in savings to put down towards the purchase price and plans to borrow the rest. She is interested only in 30-year, fixed-rate mortgages and from her research on bankrate.com, Luna has assembled the following list of mortgage products from which to choose.

Mortgage	APR (%)	Points (%)
A	5.75	0.00
B	6.00	-0.50
C	5.60	0.10
D	5.65	-0.20
E	5.50	1.50

The APR is compounded monthly. Points correspond to a fee (or rebate when negative) paid at origination. The fee is computed as the product of the points and the principal amount of the loan. Luna's annual opportunity cost is 5.904%, and you may assume that Luna plans on staying in the home for at least 30 years unless otherwise stated.

Using this information, answer the following questions.

- a. Can Luna determine the best option by comparing the APRs of each loan and selecting the lowest one? (Yes or No and briefly explain.)
- b. What is the monthly payment corresponding to each loan?
- c. What is the NPV of each loan?
- d. What is the IRR of each loan? How does each compare to the corresponding APR? Explain the relation.
- e. Based on your answers to the previous questions, which loan should Luna select? Briefly explain your answer.
- f. How much does Luna owe on each loan after five years (i.e., immediately after her 60th mortgage payment)?
- g. Recompute the NPV of each loan option assuming that Luna remains in the home for five years, selling her home for \$1 million just after making her 60th mortgage payment. Which loan should Luna select? If it differs from your answer to the previous question, briefly explain why.

6.17 Bryn Mawr college completed construction several years ago on a new academic building that cost \$3.6 million. At that time, the college decided to lease the space to a local business for \$625,000 per year. The business has five years left on their lease agreement, and the next lease payment is due one year from today. Bryn Mawr is deciding whether to terminate the lease agreement, and modify and use the space for its own purposes.

The modifications require \$500,000 of capital expenditures today, which would straight-line depreciate to zero over the next five years. The building would allow for an increase in the number of students and corresponding increase in revenue. One year from today, revenue is forecast to increase by \$850,000, a figure that will grow by 3% per year. Building maintenance - operating expense - is \$195,000 per year and will also grow at 3% per year.

Because the college is a nonprofit entity, they pay no taxes on income generated from educating their students. However, they do pay taxes, at a rate of 21%, on income from non-educational activities, notably the income from the building lease.

The college's cost of capital is 8% per annum, and their choice between continuing the lease and modifying and using the building has no effect on what happens after five years from today.

Using this information, answer the following questions.

- a. What is the value today of the remaining after-tax lease payments?
- b. What are Bryn Mawr's revenue forecasts for its use of the building?
- c. What are Bryn Mawr's EBITDA forecasts for its use of the building?
- d. What are Bryn Mawr's EBIT forecasts for its use of the building?
- e. What are Bryn Mawr's NOPAT forecasts for its use of the building? How do they compare to the EBIT forecasts? Explain.
- f. What are Bryn Mawr's free cash flow forecasts for its use of the building? How do they related to the EBITDA forecasts? Explain.
- g. What are the NPV and IRR of using the building?
- h. Should Bryn Mawr continue with the lease or modify and use the building?
- i. At what tax rate would Bryn Mawr be indifferent between its two options - continuing the lease and modifying and using the building?

Chapter 7

Investing: Fixed Income

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

This chapter examines fixed income investing. **Fixed income** refers to investments in which there is a fixed schedule of payments, if not a fixed amount for each payment. We've already examined several fixed income investments from the perspective of the borrower including mortgages, leases, student loans, auto loans, and credit cards. Now we take the lender's or investor's perspective.

More precisely, this chapter explores **savings accounts**, **certificates of deposits (CDs)**, and **bonds**. All of these investments are simply loans in which we, the investor, lend money to an entity such as a bank, government, or corporation in exchange for the return of our money plus some interest in the future. Specifically, this chapter

- examines common savings vehicles offered by banks, such as savings accounts and certificates of deposit,
- introduces bonds as an investment for savers and source of funding for governments and corporations,
- shows that bond valuation is a straightforward application of the Fundamental Value Relation,

- defines a bond yield and illustrates the yield curve,
- distinguishes between expected returns, the r in our Fundamental Value Relation and what we *expect* to earn in the future as investors, and realized returns, what we *actually* earn.
- considers the implications of taxes on investor returns,
- identifies and quantifies the risks associated with investing in bonds,
- applies our Fundamental Value Relation to answer several questions including:
 - How can we take advantage of pricing errors in bond markets?
 - How can we hedge, or mitigate, the interest rate risk of our bond portfolio?
 - How does inflation risk affect our real bond returns, and how can we hedge that risk?
 - How does default risk impact our bond returns?

7.1 Savings Accounts and Certificates of Deposit

Perhaps the most natural place to save our money is in a bank or credit union.¹ The two most popular savings products that these financial institutions offer are **savings accounts** and **certificates of deposits** or **CDs**.

7.1.1 Savings Accounts

There are two types of savings accounts - traditional and high yield - both of which operate in a similar manner. We can deposit money into and withdraw money from the account at any time and at no cost. As such, money in a savings account, like a checking account, is referred to as a **demand deposit**. Unlike a checking account, savings accounts often limit the *number* of times we can withdraw from the account without a fee. Six times per month is a common limitation. While in the account, money earns interest that is compounded daily to coincide with our ability to deposit or withdraw at any time. However, *the interest*

¹From the customers perspective, banks and credit unions are fairly similar in terms of the services they offer. However, banks are for-profit entities owned by shareholders. Credit unions are not-for-profit cooperatives owned by members who share some common bond such as their industry of employment.

rate on a savings account can change at any time. This last point is important for investors because what may seem like a very attractive interest rate on our savings may not last long.

Figure 7.1 shows the interest rates on a traditional savings accounts at Chase bank as of May 19, 2023. There are several points worth noting. Two different rates are quoted - “Interest Rate” and “APY.” The Interest Rate is the APR and measures the simple interest earned in a year. The APY, or annual percentage yield, is the EAR which accounts for daily compounded interest.

Chase Premier Savings SM									
Earn Premier relationship rates when you link the account to a Chase Premier Plus Checking or Chase Sapphire Checking account, and make at least five customer-initiated transactions in a monthly statement period using your linked checking account.									
RELATIONSHIP RATES			STANDARD RATES		RELATIONSHIP RATES			STANDARD RATES	
Balance	Interest Rate	APY	Interest Rate	APY	Balance	Interest Rate	APY	Interest Rate	APY
\$0-\$9,999	0.02%	0.02%	0.01%	0.01%	\$50,000-\$99,999	0.02%	0.02%	0.01%	0.01%
\$10,000-\$24,999	0.02%	0.02%	0.01%	0.01%	\$100,000-\$249,999	0.02%	0.02%	0.01%	0.01%
\$25,000-\$49,999	0.02%	0.02%	0.01%	0.01%	\$250,000+	0.02%	0.02%	0.01%	0.01%

Figure 7.1: Checking and Savings Rates from Chase Bank (May 19, 2023)

While the rates are similar, there are many different quotes that vary by how much money is in the savings account (i.e., “Balance” in the figure) and whether or not the account is linked to a checking account that is actively used (Standard vs. Relationship rates). Chase, like most other banks, also offers different rates depending on the tier of service, which are not shown in the figure. In general, more money with the bank means higher interest rates on savings products, 0.02% versus 0.01% in this case. The higher interest rate is the bank’s way of incentivizing us to give the bank more of our money and utilize more of its services, both of which add up to more money for the bank.

Determining how much money we will have in our savings account at any date in the future is difficult because the interest rate can and does change over time. Assuming the rate does not change over the next year, investing \$10,000 in a Chase savings account for one year at an EAR (APY in bank lingo) of 0.02% is just a matter of computing a future value.

$$10,000 \times (1 + 0.0002) = \$10,002.00$$

After one year, we would earn \$2. Not very appealing.

Contrast this result with what we would have if we invest in a high yield savings account. For example, Marcus, the online bank owned by Goldman Sachs, was offering a 4.15% APY

(i.e., EAR) in May of 2023. Assuming the rate remains unchanged, a one year investment of \$10,000 would generate

$$10,000 \times (1 + 0.0415) = \$10,415.00.$$

Like choosing a mortgage, it pays to shop for a savings account.

Another important feature of bank savings accounts is that as long as the balance in our account is below a certain threshold, we are *guaranteed* to receive our money should the bank go bankrupt. The Federal Deposit Insurance Corporation (FDIC) collects fees from member banks to provide insurance on all checking and savings accounts, as well as CDs, up to \$250,000 as of 2023. The National Credit Union Administration (NCUA) provides similar insurance for deposits at credit unions. Note, this insurance covers \$250,000 across all accounts at any one bank that are owned by the same individual or company. Any money above this threshold at one bank is uninsured and at risk of loss if the bank goes bankrupt.² This risk is an example of **counterparty risk** in which the other side of a transaction - the bank in this example - fails to deliver on its promises.

When a bank fails, and many do, the FDIC steps in to take over the bank and ensure a seamless transition from the customer's perspective. Figure 7.2 shows the number of bank failures by year in the United States and the total assets of the banks that failed. Importantly, not a single insured dollar was lost by a customer during this time. So, as long as the sum of our balances are below the insurance cap, there is no need to worry about our money in the bank.

7.1.2 Certificates of Deposit (CDs)

We deposit money into a Certificate of Deposit or **CD**, which like a savings account is FDIC insured. However, CDs differ from savings accounts in several important ways. First, once we open a CD by making an initial deposit, we typically can't add money to the account in the future. Some banks offer **add-on CDs** enabling us to contribute more money; other banks with whom we have a relationship may allow us to transfer money from another account by paying a fee.

Second, we can't withdraw principal, though we often can withdraw earned interest, from the CD before the end of the term. Opening a CD means committing to leave our money in

²The failure of Silicon Valley Bank in 2023 was an exception to this rule. The federal government guaranteed *all* deposits, including money in excess of the \$250,000 limit, to avoid the spread of a bank run to other banks.




Figure 7.2: Bank Failures - 2001-2023 (Federal Deposit Insurance Company)

the account for a certain amount of time - anywhere from one month to 10 years. As such, CDs are referred to as **time deposits**. Figure 7.3 presents rate quotes for different CDs offered by Chase bank. CD rates vary depending on how much we invest, for how long we invest (i.e., the CD term), and whether or not we have another relationship with the bank such as a checking account.

If we do withdraw principal from the CD before the end of the term, we'll face a penalty that depends on the term of the CD. For Chase, the penalty is a loss of interest, the details of which are presented in table 1.

CD Term (Months)	Penalty (Days of interest lost)
< 6	90
6 to 23	180
> 23	365

Table 1: Early Withdrawal Penalties from Chase CDs (May 2023)

Finally, CDs guarantee a fixed interest rate over their entire term, which makes financial planning easier. We know exactly how much money we will have at the end of the CD's term. For example, a \$200,000 deposit in a 12-month CD earns a relationship rate of 3.75% according to figure 7.3. At the end of 12 months, assuming we do not withdraw any money along the way, our CD will have an ending balance of

$$200,000 \times (1 + 0.0375) = \$207,500.$$

Featured Terms (Months)	CD RELATIONSHIP RATES						CD STANDARD RATES	
	\$1,000-\$9,999		\$10,000-\$99,999		\$100,000+		\$1,000+	
	Interest Rate	APY	Interest Rate	APY	Interest Rate	APY	Interest Rate	APY
3 - 5	1.98%	2.00%	3.44%	3.50%	3.92%	4.00%	0.01%	0.01%
12 - 14	2.96%	3.00%	3.20%	3.25%	3.68%	3.75%	0.01%	0.01%
24 - 29	1.98%	2.00%	1.98%	2.00%	1.98%	2.00%	0.01%	0.01%
Other Terms (Months)¹								
1	0.02%	0.02%	0.02%	0.02%	0.02%	0.02%	0.01%	0.01%
2	0.02%	0.02%	0.02%	0.02%	0.02%	0.02%	0.01%	0.01%
6 - 8	2.96%	3.00%	2.96%	3.00%	2.96%	3.00%	0.01%	0.01%
9 - 11	0.02%	0.02%	0.05%	0.05%	0.05%	0.05%	0.01%	0.01%
15 - 17	0.02%	0.02%	0.05%	0.05%	0.05%	0.05%	0.01%	0.01%
18 - 20	0.02%	0.02%	0.05%	0.05%	0.05%	0.05%	0.01%	0.01%
21 - 23	0.02%	0.02%	0.05%	0.05%	0.05%	0.05%	0.01%	0.01%
30 - 35	1.49%	1.50%	1.49%	1.50%	1.49%	1.50%	0.01%	0.01%
36 - 41	1.98%	2.00%	1.98%	2.00%	1.98%	2.00%	0.01%	0.01%
42 - 47	1.49%	1.50%	1.49%	1.50%	1.49%	1.50%	0.01%	0.01%
48 - 120	1.49%	1.50%	1.49%	1.50%	1.49%	1.50%	0.01%	0.01%

Figure 7.3: Certificate of Deposit (CD) Rates from Chase Bank (May 19, 2023)

When the CD matures, we want to withdraw all of the money because it will no longer earn interest unless we invest in non-traditional CDs.³

7.2 What are Bonds?

Savings accounts and CDs are **nonmarketable** or **nontradeable**, meaning we can't buy or sell a savings account or CD to someone else. Once we open a savings account or CD, it's ours until we close it. And, if we need to withdraw money from a CD before the end of its term, we face a stiff penalty.

Bonds, like savings accounts and CDs, are simply loans. Governments (federal, state, local, foreign), corporations, federal agencies, financial institutions, and educational institutions all borrow money from investors by issuing bonds. However, unlike savings accounts and CDs, most bonds are **marketable** or **tradeable**.⁴ They can be bought and sold by different investors multiple times over their lives. So, while some investors buy and hold a bond until it matures, most don't.

³An **automatically renewable CD** takes the money from an existing CD - interest and principal - and opens a new CD. A **CD ladder** consists of several CDs of different terms. For example, a 12-month ladder may consist of a 3, 6, 9, and 12 month CD. As each CD matures, a new 12-month CD is opened with the proceeds of the maturing CD.

⁴Series EE and I bonds issued by the U.S. government are examples of bonds that are not marketable.

7.2.1 Maturity or Term

Bonds differ in several regards beginning with the **maturity** or **term**, which determines the length of the loan. Some bonds - bought and sold by financial institutions - mature in less than a day. Other bonds never mature - so-called perpetuities. Their terms are infinite. Yale University owned such a perpetual bond issued in 1648 by a Dutch company that continued to pay interest through 2023. Most bonds have maturities ranging from one month to 30 years.

7.2.2 Repayment Schedules

Bonds also differ in their repayment terms, or how the borrower must repay the loan. Bonds that require interest be paid on scheduled dates over the life of the loan and the principal at maturity are called **coupon bonds**. The principal amount of the bond is also referred to as the **par** or **face** value of the bond. Figure 7.4 illustrates this repayment scheme. The name “coupon” comes from how interest was collected in the 19th and early 20th century. Bond owners received a booklet in which each page corresponded to a different payment that the bond owner was owed. To receive payment, bondholders would exchange each page in the booklet for the corresponding payment. The pages corresponding to the interest payments were referred to as coupons.




Figure 7.4: Repayment Scheme for Coupon Bonds

Some bonds make no explicit interest payments, only the principal at maturity as illustrated in figure 7.5. As such, these bonds are called **zero coupon bonds** or just **zeros**. Interest is earned on these bonds by purchasing them at a discount to the principal amount. As such, these bonds are sometimes referred to as **discount bonds**.

Both coupon and zero coupon bonds are referred to as **bullet** bonds because of the large principal payment at maturity. In contrast, **amortizing** bonds pay bond owners interest and




Figure 7.5: Repayment Scheme for Zero Coupon Bonds

principal over the life of the loan. We saw examples of amortizing loans, such as mortgages, student loans, and auto loans in chapter ??.

7.2.3 Security and Priority

Some bonds are **secured** by the asset's of the borrower meaning that if the borrower fails to repay their debt, the lender (i.e., bond owner) can seize the assets securing the debt. These assets act as **collateral** for the lender, which can reduce the interest rate charged on the loan. For example, mortgages are amortizing loans secured by homes. Auto loans are amortizing loans secured by cars. Corporations often promise or pledge some or all of their assets as collateral for their loans.

Security is closely related to a bond's **seniority** or **priority** when a borrower defaults. Many companies issue multiple bonds. Priority determines the order of payment in default - who gets paid first, second, etc. - when courts adhere to the **absolute priority rule**, which unfortunately has the same acronym **APR** as the annual percentage rate. The list below provides a sample debt priority structure for a corporation where the highest priority is listed at the top.

1. Senior secured (first-lien)
2. Secured (second-lien)
3. Senior unsecured
4. Senior subordinated
5. Subordinated
6. Junior subordinated

7.2.4 Covenants

Another important feature of bonds, and loans more broadly, is **covenants**, which are contract provisions specifying the expectations of both parties to the loan. While the behavior of both parties is important, covenants affecting the borrower tend to be more extensive and onerous because it is the lender who typically has the most to lose, namely, their money. There are several ways to categorize covenants. One useful categorization is negative and affirmative. As their names suggest, negative covenants detail actions the borrower is not allowed to undertake, affirmative covenants detail the actions the borrower must take.

Negative covenants can limit or prohibit a variety of borrower activity including:

- Additional borrowing or leasing activity,
- Payments to shareholders via dividends or stock buybacks,
- Asset sales,
- Mergers and acquisitions,
- Pledging assets as collateral for other debt, referred to as a **negative pledge clause**,
- Letting financial ratios exceed or fall short of different thresholds such as a maximum leverage (debt-to-ebitda) ratio or minimum interest coverage (ebitda-to-interest) ratio.
- Capital expenditures or investments in financial assets, and
- Changes in the ownership of the company or the composition of the board of directors, sometimes referred to as a **posion put**.

Affirmative covenants require borrowers to:

- Use the bond proceeds for specific purposes,
- Pay interest and principal in a timely manner,
- Maintain accurate and accessible financial records,
- Comply with all relevant laws,
- Maintain all assets in good working order,
- Purchase insurance for properties and equipment,

- Allow lenders to inspect the companies financial records and assets, and
- Notify lenders of any event of default.

There are two types of default. A **payment default** occurs when the borrower fails to make a required interest or principal payment. A **technical default** occurs when any other covenant violation occurs. When a default occurs, lenders can accelerate the debt, that is require borrowers to repay all of the debt immediately. Alternatively, creditors can renegotiate the debt with borrowers in an effort to avoid bankruptcy. For bonds with many dispersed investors, renegotiation is relatively uncommon because of the cost. There are **free-rider problems** in which a small number of lenders bear all of the costs of renegotiating and **collective action problems** in which it is difficult to get a large number of creditors to agree to a renegotiation. In contrast, bank loans for which ownership of the debt is concentrated among few investors is quite often renegotiated to avoid costly bankruptcies.

7.2.5 The Most Important Bonds? U.S. Treasury Securities

While many entities issue bonds, the largest and arguably most important issuer in the world is the U.S. federal government. The bonds issued by the U.S. government are referred to as **Treasury securities** or **Treasurys**. As of June 2024, there were \$35 trillion of Treasury's outstanding. In other words, the U.S. federal government owes its bond investors (i.e., creditors) \$35 trillion.⁵ The Treasury market is also the most **liquid** bond market with \$880 billion of securities traded in May 2024 alone. As a quick aside, liquidity of a market or asset refers to the ease with which one can transact. Assets that can be bought and sold quickly and cheaply, like publicly listed stocks or Treasury bonds, are examples of liquid assets. Real estate, venture capital, and private equity investments are examples of illiquid assets.

In addition to their size and liquidity, Treasury bonds are viewed as a safe, if not the safest, investment available to most investors. As such, they are often used as a proxy for a **risk-free asset** because investors believe they are guaranteed by the US government to receive all of their promised interest and principal.⁶ As such, the interest rate on Treasury securities is typically lower than that on all other types of loans in the U.S. However, as we'll see, they are only risk-free in a narrow sense.

⁵The plural of Treasury when referring specifically to Treasury bonds bonds spelled as "Treasurys."

⁶The U.S. government has defaulted several times in its history; first on on demand notes in 1862 and then on gold bonds in 1933. It also failed to redeem its silver certificate paper dollars for silver dollars in 1968, and its dollars held by foreign governments for gold in 1971.

The reason the U.S. government borrows so much money by issuing Treasury securities is its excessive spending. The government spends money on defense, social security, healthcare programs (e.g., medicare and medicaid), education, infrastructure, etc. It pays for these things with the tax and fee revenue it raises from people and businesses. However, most of the time, the government spends more than it raises as seen in figure 7.6. The figure presents the quarterly budget surplus (when positive) or deficit (when negative) for the U.S. federal government. Most of the time the line is below the horizontal axis implying that the government spent more money than it raised. That excess spending was extreme in the wake of the financial crisis of 2008 and the Covid pandemic in 2020 and 2021. As a result, the Treasury department holds auctions throughout the year in which it sells bonds to investors to pay for the excess spending.




Figure 7.6: Quarterly U.S. Federal Surplus or Deficit

There are three different types of Treasury securities differentiated by their maturities and in some cases their repayment schemes.

1. **Treasury bills** or **T-bills** are short-term, zero-coupon bonds that come in five different maturities: 4-, 8-, 13-, 26-, and 52-weeks.
2. **Treasury notes** or **T-notes** are medium-term, semi-annual coupon bonds that come

in five different maturities: 2-, 3-, 5-, 7-, and 10-years.

3. **Treasury bonds** or **T-bonds** are long-term, semi-annual coupon bonds that come in 20- and 30-year maturities.

While Treasuries are referred to as bills, notes, or bonds depending on their maturity, they are all examples of “bonds,” which are ultimately just loans to the US government for different lengths of time.

Finally, though the Treasury doesn’t issue any zero coupon bonds with a maturity greater than one year, some financial institutions create such bonds by **stripping** the coupons and principal from T-notes and T-bonds and selling them as new securities. This process is illustrated in figure 7.7 in which a 30-year Treasury bond is stripped into 61 STRIPS. Sixty of the STRIPS come from the coupons and are referred to as **interest STRIPS** or **C-STRIPS**. The principal from the Treasury note makes up **principal STRIPS** or **P-STRIPS**.⁷

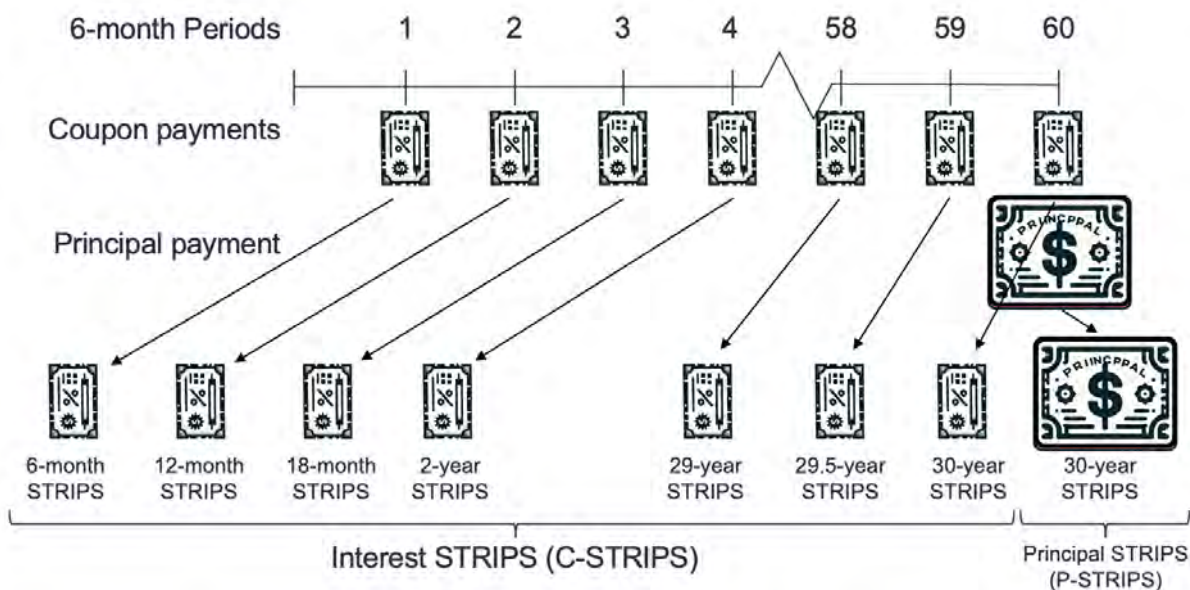


Figure 7.7: Creating Treasury STRIPS from a 30-year Treasury Bond

For the remainder of this chapter, we’ll rely on Treasuries to illustrate different concepts. In addition to being practical and relevant, doing so avoids some of the complications that arise when analyzing other types of bonds with uncertain cash flows. As we mentioned, above, we can treat Treasury bonds as if their are default-free; the U.S. federal government

⁷Though economically very similar, some foreign governments tax C-STRIPS and P-STRIPS differently, which can lead to price differential across otherwise identical securities.

will always pay its debt (...hopefully). In addition to their safety, there are no **embedded options** in Treasury bonds. Treasuries provide neither the borrower nor the lender the option of early retirement of the bonds. That is, the Treasury cannot force investors to sell their bonds back to the Treasury prior to maturity, and investors cannot demand the Treasury repay bond principal prior to the maturity.⁸ However, investors are free to buy and sell bonds to other investors at any time during the life of the bond, hence Treasuries status as marketable investments. Because of this certainty, the cash flows of Treasury bills, notes, bonds, and STRIPS are all perfectly predictable. Investors know exactly how much money they will receive and when they will receive it.

Towards the end of the chapter and in the technical appendix, we will briefly explore bonds issued by corporations that may default and fail to make their promised payments. We'll also explore **Treasury Inflation Protected Securities** or **TIPS** whose cash flows depend on future inflation, which is not perfectly predictable.

7.3 Investing in T-bills and STRIPS

7.3.1 Valuation

Valuing zero coupon bonds such as T-bills and STRIPS is simply a matter of finding the present value of one future cash flow - the face value of the bond. For example, let's find the value of a Treasury STRIP with a \$1,000 par value and five years to maturity. Convention in the Treasury market is to compound interest semi-annually for consistency with the payment frequency of the coupon bonds - T-notes and T-bonds. So, let's assume a semi-annually compounded APR of 10%. The timeline for this bond is presented in figure 7.8, which shows time measured in both six-month periods and years.




Figure 7.8: Timeline for a 5-Year Treasury STRIP

⁸An embedded option gives the borrower or lender the option to alter the maturity of the bond. For example, **callable bonds** include a call provision that allows the bond issuer to purchase or retire the bond at certain times during the bond term for a pre-specified price. **Puttable bonds** include a put provision that allows the bond investor to force the borrower to retire the bond at certain points during the bond term usually at par. In other words, the bond issuer will have to buy back all the bonds at their par value.

For example, the price of \$12,800,000 of par value for the STRIP we just examined can be computed by multiplying the total par value by the price per \$1 of par value like so.

$$\text{\$12,800,000 par value} \times \underbrace{\frac{\text{\$61.39 market value}}{\text{\$100 par value}}}_{\text{Price per \$1 of par value}} = \text{\$7,858,089.65.}$$

Alternatively, we can multiply the number of bonds by the price per bond.

$$\underbrace{\frac{\text{\$12,800,000 total par value}}{\text{\$100 par value per bond}}}_{\text{Number of bonds}} \times \text{\$61.39 price per bond} = \text{\$7,858,089.65.}$$

Both interpretations produce the same result.

Second, the value of any zero coupon bond, including STRIPS and T-bills, is always less than the face value as long as the interest rate is positive.¹⁰ This is why zero coupon bonds are referred to as discount bonds or sometimes **pure discount bonds**. The difference between the price and par value of a zero coupon bond corresponds to the interest.

Finally, we've been computing bond values, which need not equal bond prices at which we can buy and sell bonds. When prices deviate from value is when we have an opportunity to profit from our investments beyond what we should earn for the risk we are taking. And, while we might think we're not taking any risk when we buy a Treasury security because the government is very unlikely to default, we in fact are taking a risk that we'll discuss below.

Let's do one more example, this time valuing a 13-week T-bill that we want to buy on May 18, 2023 and that matures on August 17, 2023. The timeline is in figure 7.10, which shows one cash flow of \$100, 91 days in the future. Assuming the semi-annual compounded APR is 5.2451%, and therefore the EAR is $(1 + 0.052451/2)^{1/2} - 1 = 5.3139\%$, the T-bill's value can be estimated like so.

$$Value_0^{13-week} = \frac{100}{(1 + 0.052451/2)^{91/182.5}} = \frac{100}{(1 + 0.053139)^{91/365}} = \$98.72$$

7.3.2 Prices and Yields

We can buy and sell bonds on most brokerage platforms or through the government at TreasuryDirect. In most cases, we'll be presented with price quotes such as those in Table

¹⁰Nominal interest rates can and have gone negative, though it is relatively rare. Interest rates on government debt in Switzerland, Denmark, Sweden, and the Euro area were all negative at some point in 2016. Negative interest rates reflect the price investors are willing to pay governments, and in some cases companies, to hold their cash.




Figure 7.10: Timeline for Newly Issued T-Bills

2. The first bond is a 1-year T-bill issued on May 16, 2023. The second bond is a STRIP constructed from the principal of a 30-year Treasury bond originally issued on May 15, 2000. The last bond is a STRIP constructed from the coupon of a different 30-year Treasury bond issued on February 15, 2001.

Security Type	Original Term	Maturity Date	Bid/Ask Price (per \$100 par)	Bid/Ask Quantity (\$ par)	Bid/Ask YTM (%)
Bill	52-week	5/16/2024	95.23/95.25	10,000/10,000	5.18/5.16
P-STRIP	30-year	5/15/2030	77.21/77.52	5,000/3,000	3.75 /3.70
C-STRIP	29.5-year	8/15/2030	76.05/76.30	2,000/2,000	3.84/3.79

Table 2: Bond Price Quotes as of June 1, 2023 (Sources: Treasury Direct and JP Morgan)

There are two price quotes for each bond - a bid and an ask. The **bid price** is the price at which the market is willing to buy a bond from us (i.e., the price at which we can sell a bond). The **ask price** is the price at which the market is willing to sell a bond to us (i.e., the price at which we can buy a bond). Having a different price for the same bond depending on whether we want to buy or sell may seem odd, if not frustrating because the price at which we can buy the bond is higher than the price at which we can sell it. However, this difference corresponds to a **transaction cost** investors pay for the ability to buy and sell bonds quickly and at a reasonable price. Other investors called **dealers** stand ready to buy and sell bonds or provide **liquidity** to us. The price they charge for this service is reflected in the difference between the bid and ask prices known as the **bid-ask spread**.

The second to last column in table 2 shows how many bonds we can buy or sell at the current price quotes. One bond on J.P. Morgan's trading platform has a face value of \$1,000, implying we can buy or sell up to $10,000 \times 1,000 = \$10,000,000$ of *par* value worth of T-bills at the ask and bid prices, respectively. Orders larger than 10,000 bonds may not be able to be executed at the quoted prices because of what is available in the dealers inventory. Dealers hold a limited amount of each type of bond and a sufficiently large order may require them to engage with other dealers to fulfill the order.

The last column presents the **yield-to-maturity (YTM)** or **required yield** for each bond. We'll refer to this measure more succinctly as the bond **yield**.¹¹ The yield is the one discount rate such that the price of the bond equals the sum of the future cash flows of the bond discounted by that rate. Let's compute the bid yield for the P-STRIP from Table 2.

$$77.21 = \frac{100}{1 + \text{yield}}^{2,540/365.25} \implies \text{yield} = \left(\frac{100}{77.21} \right)^{365.25/2,540} - 1 = 3.7983\%$$

Computing the yield like this generates the EAR of the bond, which tells us how much we earn on an annual basis assuming we hold the bond to maturity. The periodic or semi-annual yield is computed similarly, only measuring time in 6-month periods.

$$77.21 = \frac{100}{1 + \text{yield}}^{2,540/182.625} \implies \text{yield} = \left(\frac{100}{77.21} \right)^{182.625/2,540} - 1 = 1.8770\%$$

This number tells us how much we earn each six months and is slightly less than one half the EAR because of compounding.

In practice, when we talk about a bond's yield to maturity, we typically do so using the market convention of a **bond equivalent yield**, which is bond lingo for the APR. So the bid yield would be quoted as follows.¹²

$$\text{yield} = 1.8770\% \times 2 = 3.7540\%$$

More generally, the yield to maturity for a zero coupon bond can be expressed as follows.

$$\text{Price}_0 = \frac{\text{Face}}{(1 + \text{yield}/k)^T} \implies \text{yield} = \left[\left(\frac{\text{Face}}{\text{Price}_0} \right)^{1/T} - 1 \right] \times k \quad (7.1)$$

The maturity, T , is measured in periods and k is the compounding frequency per year. Aside from the lingo, there is nothing new here. We are using the present value, which we're calling price, of a single future cash flow, which we're calling face or par value, T periods in the future to find the APR, which we're calling yield.

¹¹There are other "yields" in bond markets, such as **yield-to-worst** and **yield-to-call**, however these concepts are beyond the scope of this book. A good reference is Frank Fabozzi's book, *Bond Markets, Analysis, and Strategies*.

¹²T-bills actually have their own convention, different from other Treasuries and Corporate bonds, in which they yield is computed on a **bank discount basis**.

$$\text{yield} = \frac{\text{Par Value} - \text{Price}}{\text{Par Value}} \times \frac{360}{\text{Days to maturity}}$$

The yield to maturity is the internal rate of return (IRR) of a bond. Recall from chapter 5 that the IRR is the one discount rate such that the net present value (NPV) of an investment equals zero. By setting the present value of the future cash flows equal to the price, the yield ensures that the NPV of the bond investment is zero or that the bond is fairly priced. Mathematically, this can be seen by taking our Fundamental Value Relation in which value and price are assumed to be equal, and moving the price to the right side of the equation to get an expression for NPV.

$$Price_0 = \frac{Face}{(1 + APR/k)^T} \implies 0 = \underbrace{-Price_0 + \frac{Face}{(1 + APR/k)^T}}_{NPV}$$

The bond yield is the APR in the above expression.

Bond prices and yields contain the same information. If we know the yield on a bond, we can derive its price and vice versa. However, we can't lose site of the fact that bond prices are the fundamental element in these equations. Investors buy and sell bonds, which affect bond prices, which in turn dictate bond yields. So, when investors quote a bond yield, it's simply a different way of quoting a bond price.

Prices and yields move in the opposite direction of one another. When bond yields increase, bond prices decrease and vice versa. But, this is analogous to saying that when discount rates increase, value decreases and vice versa. The bond price-yield relation is illustrated in figure 7.11 for a 10-year Treasury STRIP with a 4.36% yield, but the slope and shape of the relation apply to all bonds without embedded options.

7.3.3 Yield Curve and the Term Structure of Interest Rates

Table 2 above shows that on June 1, 2023 bonds of different maturities have different yields. For example, the T-bill has 350 days to maturity and a bid yield of 5.18%. The C-STRIP has 2,632 days, or approximately 7.2 years, to maturity and a bid yield of 3.8%. In fact, bonds of different maturities almost always have different yields. By graphing the yields of bonds with different maturities, we produce a **yield curve**. Figure 7.12 presents the yield curve and corresponding bond prices as of September 2, 2023 for Treasury STRIPS with maturities ranging from one to 20 years.

The right and left sides of the figure are referred to as the **short end** and the **long end** of the yield curve because they reflect the yields of short- and long-term bonds, respectively. The short end of the curve in Figure 7.12 is decreasing before gently rising around seven




Figure 7.11: Price Yield Relation for 10-Year Treasury STRIP

years and then eventually decreasing at the very long end of the curve. Thus, short-term bonds have relatively high yields compared to medium- and very long-term bonds.

If we're being precise, and we should be, Figure 7.12 presents the *spot* yield curve. **Spot yields** (a.k.a., **spot interest rates** or **spot rates**) correspond to the interest rate on a loan in which the lender gives money to the borrower at or around the time of the agreement. Most loans with which people and companies engage rely on spot rates. However, parties also engage in agreements to lend money at some time in the future. The interest rate agreed upon today for a loan to be made in the future is referred to as a **forward interest rate** or **forward rate**. See the technical appendix of this chapter for more details on forward rates.

If yields are simply a reflection of prices, we might ask why not graph bond prices instead of bond yields? Looking at the prices in Figure 7.12, this plot wouldn't be terribly informative because the prices of zero coupon bonds will almost always decline with maturity. This fact follows from the Fundamental Value Relation applied to zero coupon bonds.

$$Price = \frac{Par}{(1 + r)^T}$$

As T increases, price declines all else equal. So, comparing the prices of zero coupon bonds at different maturities is not particularly insightful. For example, the lower price on the 20-year bond (\$42.12) compared to that on a three-year bond (\$87.35) is largely a result of

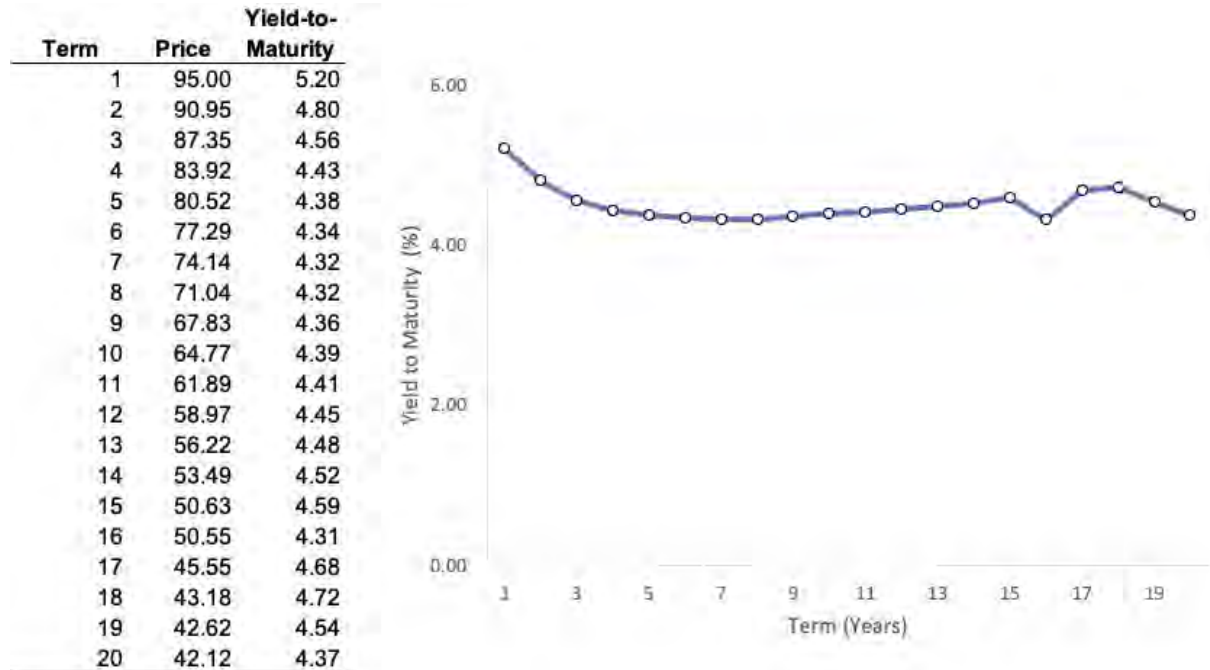


Figure 7.12: Treasury STRIP Yield Curve, September 2, 2023

having to wait longer to receive our money. It doesn't tell us which pays more interest on an periodic basis.

In contrast, we can compare the yields across different maturities to determine which are promising more or less interest on an annual basis. The yields on the 20- and 3-year bond, 4.37% and 4.56%, tell us that the 3-year offers more annual interest than the 20-year. While informative, this does not mean that the 3-year bond is necessarily a better investment than the 20-year bond. The 3-year bond only guarantees us this interest rate for three years. If interest rates decrease below 4.37% after three years, we may regret our short-term investment because the 20-year bond guarantees 4.37% for 20 years. Conversely, if we buy the 20-year bond and interest rates increase above 4.37% then we may regret having bought the longer-term bond.

In sum, the yield curve is simply a plot of the yield to maturities corresponding to bonds of different maturities. It illustrates the **term structure of interest rates**, or the relationship between loan terms and interest rates that is discussed in more detail below. The key implication of the yield curve is that the discount rate in our Fundamental Value Relation, r , can differ across future periods.

$$Value = \frac{CashFlow}{1 + r_1} + \frac{CashFlow}{(1 + r_2)^2} + \frac{CashFlow}{(1 + r_3)^3} + \dots \quad (7.2)$$

Each period has a possibly different discount rate, r_1, r_2, r_3, \dots . Using a constant discount

rate as we've done up to this point is only a simplifying assumption, but often a reasonable one. When evaluating fixed income investments, assuming a constant discount rate is no longer a reasonable assumption.

7.3.4 Liability Driven Investment: Cash Flow Matching

An important use for fixed income investments is **liability driven investment**, or **LDI**. LDI is an investment approach aimed at meeting future liabilities or expenses. LDI can be found everywhere. For example, we might structure our personal investments to meet future tax obligations, education expenses, retirement consumption, or to fund a large purchase such as a car or home. Similarly, businesses structure the investment of their money to fund future expenses, capital expenditures, and acquisitions. Insurance companies invest their money to ensure they can meet future claims by policy holders. Defined benefit pension plans - organizations tasked with investing employee retirement contributions - structure their investments to meet the future benefits paid to retirees.

In some sense, *all* investment is liability driven. We - people and businesses - save and invest money to spend in the future. That said, the LDI label is typically reserved for investments needed for explicit future liabilities and expenses, as opposed to consumption goals such as a car or house. So, let's illustrate LDI with an example.

The Problem

A law firm - Pearson, Specter, and Litt (PSL) - is **acquiring** (i.e., buying) an information technology firm - Benji Inc. The deal is expected to close six months from today on August 15, 2024. At the deal close PSL - the **acquirer** - will pay \$40 million to the owners of Benji - the **target**. PSL will finance the acquisition with \$15 million of its own money (i.e., cash) and \$25 million of debt from their bank. This transaction is an example of a business **acquisition**.

The questions in which we're interested are: How should PSL store its cash for the next six months, and how much money does PSL need today to ensure it has enough cash when the deal closes in six months on August 15, 2024? In the context of LDI, the liability is the \$40 million PSL is facing to purchase Benji in six months. The investment today to meet that liability is what we are trying to determine.

Investment Options

Let's look at several options and weigh the risk and return characteristics of each investment. While some options may seem obviously silly, the discussion is valuable because it illustrates the importance of identifying the risk-return tradeoffs for any investment and shows how to construct an appropriate LDI strategy.

1. *Checking account.* Interest rates are typically quite low, often near or at zero. Because of this, there is little interest rate or return risk as we're not really earning anything on our money. There is counterparty risk because if the bank fails, PSL can only be assured of receiving what is insured - an amount well below the millions they will need for the acquisition. However, practically speaking, bank failures are relatively rare and money in excess of the insurance limit is often collected when a bank fails.
2. *Savings account.* The risk-reward profile of a savings account is similar to that of a checking account except for a possibly higher interest rate, especially if it is a high yield savings account.
3. *Certificate of deposit.* The return will be higher than checking and traditional savings accounts, but similar to high yield savings accounts. Because the interest rate on CDs are fixed, there is no return risk. PSL would know exactly how much they need to invest today. Like checking and savings accounts, this is a relatively safe investments in that most banks are unlikely to fail and if one should there is a reasonable chance that PSL can be made whole (i.e., get all their money back). However, should PSL need the money before six months, they would face a potentially severe penalty - loss of interest - for an early withdrawal. There is no flexibility with a CD.
4. *Stock market.* Stocks offer significantly higher average returns than any of the previous options and stock investments are highly **liquid** meaning they can be converted to cash quickly and with minimal cost. However, that higher average return comes with significant return risk. If the market perform poorly, PSL could find itself without enough money in six months.
5. *Real estate.* Like stocks, real estate offers a higher average return than the bank products. However, like stocks, that higher average return comes with greater risk that could put the deal in jeopardy if real estate performs poorly. Unlike stocks, real estate is **illiquid** being time-consuming and costly to convert to cash (i.e., sell).

6. *Money Market Account.* A money market account invests our money in very safe, short-term assets like T-bills, commercial paper (i.e., short-term corporate bonds), and CDs. Some provide check writing privileges, though depositors are often limited to six transactions per month. Like checking and savings accounts and CDs, money market accounts are FDIC insured. However, the interest rate on a money market account is variable and often changes over time.
7. *T-bills.* The return will be competitive with high yield savings accounts and CDs. Like CDs the interest rate on a T-bill is fixed so there is no return risk as long as the T-bill is held to maturity. PSL would know exactly how much they need to invest today. T-bills are relatively liquid meaning should PSL need the money before six months, it would be easy to sell their bond though they could lose a small amount of money. Unlike bank products, T-bills are backed by the full faith and credit of the US federal government, which is significantly safer than any bank.

Several options above are viable investments to for the PSL's purposes, specifically a high yield savings account or money market account. However, T-bills arguably offer the most attractive investment option for PSL's purpose. They ensure a market return from the safest counterparty on the planet - at least as of 2023. They lock in an interest rate over the investment horizon so PSL will know exactly what they are receiving in six months. And, T-bills preserve the flexibility to sell prior to six months should they need the cash sooner.

Investment Strategy

The primary goal of our investment strategy should be to ensure we have enough cash on hand for the deal close, while considering investment return as a secondary objective. The strategy will implement is called **dedication** or **cash flow matching**. Table 3 presents ask price quotes for several T-bills. The T-bill maturing on August 8 offers the closest maturity date to the deal close. Earlier maturity dates lose out on potential interest and later ones come with risk that we'll discuss below.

Settling on the 8/8/24 T-bills, the first question is: How many bonds should PSL purchase? They'll need \$15,000,000 in cash for the deal close. So, they should purchase \$15,000,000 of par value because the T-bills will pay their par value at maturity. Each bond has a par value of \$1,000, which means PSL must purchase

$$\$15,000,000 \text{ par value} \times \frac{1 \text{ bond}}{\$1,000 \text{ par value}} = 15,000 \text{ bonds.}$$

Security Type	Maturity Date	Ask Price (per \$100 par)	Ask Quantity (\$ par)	Ask YTM (%)
T-Bill	7/25/24	97.75	14,500	5.24
T-Bill	8/1/24	97.65	19,000	5.25
T-Bill	8/8/24	97.60	15,000	5.15
T-Bill	9/5/24	97.30	14,500	5.00
T-Bill	10/3/24	96.96	5,000	4.95

Table 3: T-Bill Ask Price Quotes as of February 15, 2024

The next question is: How much will it cost to buy these 15,000 bonds? This amount can be found by multiplying the price per bond by the number of bonds. However, because prices are quoted per \$100 of par value and a single bond has \$1,000 of par value, we need to multiply the quoted prices by 10 to get the price per bond.

$$\frac{\$97.60 \text{ market value}}{\$100 \text{ par value}} \times \frac{\$1,000 \text{ par value}}{1 \text{ bond}} = \$976.00 \text{ per bond}$$

The total cost of the 15,000 bonds is therefore

$$\$976.00 \times 15,000 = \$14,640,000.$$

To sum up, PSL only needs invest \$14,640,000 to ensure it has enough cash - \$15,000,000 - six months from now when the deal closes on August 15, 2024. It will earn

$$\left(\frac{15,000,000}{14,640,000} \right) - 1 = 2.46\%$$

over the 175 days between February 15, 2024 and August 8, 2024. This corresponds to an annualized return (EAR) of

$$(1 + 0.0246)^{366/175} - 1 = 5.21\%.$$

7.3.5 Price Dynamics and Rate Risk

Our discussion thus far has focused on buying a bond and holding it to maturity. However, as we mentioned, many investors don't buy and hold bonds until maturity. Rather, bonds are bought and sold throughout their life. This raises the questions: how and why does the value of a bond change over time?

According to our Fundamental Value Relation, value changes when (i) cash flows change, (ii) discount rates change, or (iii) the time to receiving the cash flows changes. For Treasuries,

the cash flows don't change because we assume that the U.S. will always pay its debt. What does change over time are interest rates (i.e., discount rates) and the time to receiving the cash flows. Let's illustrate the effect of these changes on bond prices using a Treasury STRIP as of February 15, 2018 with exactly five years to maturity and a current yield to maturity of 2.65%.

Figure 7.13 presents two plots. The top plot shows the evolution of the bond yield - blue series - from its issuance, February 15, 2018, to its maturity, February 15, 2023. The red series presents the yield at issuance, 2.65%. The bottom plot also presents two series. The blue series is the daily bond price. The red series is the hypothetical bond price if we assume that interest rates never change and the yield of the bond remains fixed at its initial value of 2.65%. Let's understand the mechanics behind the figures and the lessons they convey.



Figure 7.13: Interest Rate and Bond Price Evolution for 5-Year Treasury STRIP

The price of the STRIP at any point in time is just the present value of the face value,

\$100. At a 2.56% interest rate, the price per \$100 of par value is

$$Price_0 = \frac{100}{(1 + 0.0265/2)^{10}} = \$87.67.$$

If interest rates don't change, then the price of the bond at any point in time is perfectly predictable, as suggested by the straight red line. The price will steadily increase to the face value at maturity because we are getting closer to receiving our money. Mathematically, the price of the bond at any time t under the constant 2.65% interest rate assumption is given by

$$Price_t = \frac{100}{(1 + 0.0265/2)^{10-t}},$$

where t measures the number of six-month periods that have elapsed.

For example, on June 28, 2019 there are 1,328 days remaining to maturity, or $1,328/182.75 = 7.27$ six month periods. If the discount rate didn't change from 2.65%, then the bond price is

$$Price_{7.27} = \frac{100}{(1 + 0.0265/2)^{7.27}} = \$90.88,$$

which is indicated in the bottom plot by the black circle on the red series over June 28, 2019. However, interest rates *did* change between February 2018 and June 2019 and by quite a bit. On June 28, 2019, the interest rate for a loan lasting 1,328 days was 1.76% as indicated in the top plot, implying an actual price of

$$Price_{7.27} = \frac{100}{(1 + 0.0176/2)^{7.27}} = \$93.83.$$

This price is indicated in the bottom plot by the black circle on the blue series over June 28, 2019. The decline in interest rates since the issuance of the bond has led to a price increase that is greater than what would expect from just the passage of time.

In contrast, on October 19, 2018, there were 1,580 days remaining on to maturity, or $1,527/182.75=8.65$ six month periods. At the original interest rate of 2.65%, the price of our bond would have been

$$Price_{8.65} = \frac{100}{(1 + 0.0265/2)^{8.65}} = \$89.24.$$

However, the interest rate on loans of 1,580 days increased to 3.05%, implying an actual price of

$$Price_{8.65} = \frac{100}{(1 + 0.0305/2)^{8.65}} = \$87.73.$$

The increase in interest rates since the bond issuance led to a price decrease.

If we buy and hold a Treasury security until maturity, there is no interest rate risk. We know exactly how much money we will receive and when we will receive it. We can see this certainty in Figure 7.13 at the maturity of the bond where the blue and red series meet at \$100. Intersections between the two lines *before* maturity are coincidental and occur only when the interest rate equals the rate at the time the bond was issued. However, if we ever want to sell a Treasury before maturity, we face significant risk because of changing interest rates. When interest rates change, the price of the bond changes and, as we'll see next, so does our return as an investor. Interest rate risk is the primary risk facing bond investors.

7.3.6 Holding Period Returns

Continuing with the example from above, let's see how we would fare buying and selling the 5-year Treasury STRIP at different times and over different horizons. In other words, what are the return implications of trading this bond? Figure 7.14 replicates figure 7.13 and attaches to the bottom the corresponding **realized returns** over different time horizons, sometimes referred to as **holding period returns**.

For zero coupon bonds, holding period returns are computed as the ratio of money received to money spent minus one.

$$r_{t,t+s} = \frac{\text{Money Received}_{t+s}}{\text{Money Spent}_t} - 1 \quad (7.3)$$

The subscripts denote *when* we receive or spend money. So, a return, r , is measured over some time interval from t to $t + s$, the length of which is just s . These returns convey the money earned on a one dollar investment. We can measure time t and s in years or periods as long as we're consistent.

The definition in equation 7.3 is sometimes referred to as a **net return** because it measures how much money we earn excluding our initial investment. If we don't subtract one, we get a measure of **gross return** that includes our initial investment. Net returns are numbers like -0.05 and 0.10. Gross returns are numbers like 0.95 and 1.10. When people talk about returns, they're typically talking about net returns so we'll drop the "net" qualifier in our discussion.

For example, if we bought the five-year STRIP on February 15, 2018 for \$87.67 and then sold it on June 28, 2019 for \$93.83, our return over this period would be

$$\frac{93.83}{87.67} - 1 = 7.02\%$$




Figure 7.14: Interest Rate and Bond Price Evolution for 5-Year Treasury STRIP

as indicated by the blue line in the bottom graph of figure 7.14. For each dollar we invested in the bond, we earned \$0.0702. Had interest rates not changed and the price of the bond on June 28, 2019 been \$90.88 our return would have been 3.65% as indicated by the red line. The decline in interest rates to 1.76% increased the bond price, which increased our return relative to what would have happened had interest rates not changed.

If we had sold our bond on October 19, 2018, then the bond price would have been \$87.73, implying a return of

$$\frac{87.73}{87.76} - 1 = 0.07\%,$$

as indicated on the blue series in figure 7.14. Had interest rates not increased to 3.05%, the price of the bond would have been \$89.24 and our return 1.79%. Thus, changing interest rates lead to changing bond prices, which impact our investment returns.

Table 4 examines our five-year bond a little more closely. The top of the table shows the calendar dates, and corresponding six-month periods and years, on which we want to focus. The interest rates and bond prices on those dates are presented next. The top part of the bottom section of the table presents holding period returns over different horizons. For example, if we purchase the bond on February 15, 2018 and sell it one year later, our return will be

$$\frac{90.5756}{87.6665} - 1 = 3.32\%.$$

If instead we purchase the bond on February 15, 2020 and sell it two years later, our return will be

$$\frac{99.2839}{95.9439} - 1 = 3.48\%.$$

Do these results mean that the second strategy is better than the first because it generated a higher return - 3.48% versus 3.32%? No. We shouldn't compare these returns because they are measured over different time horizons. We need to standardize the time horizon of returns similar to how we needed to standardize the time horizon of recurring investments discussed in chapter 6. We do this by **annualizing** the returns, or more simply, computing the effective annual rate.

The first return, 3.32%, is measured over one year so there is nothing to do here; this is the annual realized return over this period. However, the second return, 3.48%, is measured over two years. The corresponding annual return is

$$(1 + 0.0348)^{1/2} - 1 = 1.73\%.$$

Thus, the 3.48% return earned over two years is equivalent to earning 1.73% compounded annually. Based on returns alone, the first strategy was a better investment.

Annualizing a return over a horizon *less* than a year works similarly. For example, we saw above that if we sell our bond on October 19, 2018, 246 days after purchasing it on February 15, 2018, our holding period return is 0.07%. Annualized, this return is equal to

$$(1 + 0.0007)^{365.25/246} - 1 = 0.1039\%.$$

Date (2/15)	2018	2019	2020	2021	2022	2023
Periods (6-months)	0	2	4	6	8	10
Years	0	1	2	3	4	5
Interest rate (APR %)	2.65	2.49	1.39	0.12	0.72	4.64
Price (\$)	87.6665	90.5756	95.9439	99.7604	99.2839	100.0000
Holding Period Returns (%)						
1-year		3.32	5.93	3.98	(0.48)	0.72
2-year			9.44	10.14	3.48	0.24
3-year				13.80	9.61	4.23
4-year					13.25	10.40
5-year						14.07
Annualized Returns (%)						
1-year		3.32	5.93	3.98	(0.48)	0.72
2-year			4.61	4.95	1.73	0.12
3-year				4.40	3.11	1.39
4-year					3.16	2.51
5-year						2.67

Table 4: Zero Coupon Bond Price and Annualized Return Dynamics

More generally, the annualized return of a holding period return from year t to year $t + s$ can be computed as follows.

$$\text{Annualized return} = (1 + r_{t,t+s})^{\frac{1}{s}} - 1 \quad (7.4)$$

This example illustrates several important points. First, as mentioned above, investing in bonds is risky even when the borrower is guaranteed to repay. This risk comes from changes in interest rates. Second, trading safe bonds like Treasury's is in large part betting on whether future interest rates will go up or down. If traders think rates are going down, they will want to buy bonds today. If traders think rates are going up, they will want to sell bonds today. Thus, bond traders - particularly those in the Treasury markets - are to some extent betting on interest rate movements.

Third, we can always buy and hold a bond to maturity and as long as we can be guaranteed that the borrower will repay their debt in full, we know exactly how much money we will earn. Notice that the annualized return on this strategy is 2.67%, exactly equal to the EAR of the bond when issued - $(1 + 0.0265/2)^2 - 1 = 2.67\%$. Another way to express this final implication is that no matter how many times or when a zero coupon bond is bought and sold, the compounded returns over the life of the bond must equal the holding

period return over the life of the bond. Figure 7.15 shows the holding period returns for our five-year STRIP and the two different “trading paths.”

Starting from the bottom, the first path buys the bond on February 15, 2018 and then sells on June 28, 2019 earning a return of 7.02%. The bond is bought that same day for \$93.83 and held until maturity. The return for this second leg is $100.00/93.83 - 1 = 6.58\%$. The total return over the life of the bond is $(1+0.0702) \times (1+0.0658) - 1 = 14.07\%$, or $(1 + 0.1407)^{1/5} - 1 = 2.67\%$ per year.




Figure 7.15: Interest Rate and Bond Price Evolution for 5-Year Treasury STRIP

The second path buys the bond on February 15, 2018 and then sells on September 2, 2019 earning a return of 8.70%. The bond is purchased that same day for \$95.29 and then sold again on February 28, 2020 for \$97.51, implying a return of $97.51/95.29 - 1 = 2.33\%$. The bond is purchased one last time and held until maturity earning a final return of $100/97.51 - 1 = 2.55\%$. The total return over the life of the bond is $(1+0.0870) \times (1+0.0233) \times (1+0.0255) - 1 = 14.07\%$, same as before. Thus, while there is interest rate risk to selling the zero before maturity, there is none if we hold the zero until maturity.

7.3.7 Expected Returns vs. Realized Returns

The returns computed in table 4 are **realized returns** in that they measure what has (or what we have assumed has) happened. These are different from **expected returns**, which measure what we expect but don't know will happen. In our Fundamental Value Relation, r

represents expected returns because the cash flows they are discounting occur in the future. We don't know what will happen and so we must make a guess called an expected value.

To illustrate, assume today is period 0 and we purchase the 5-year Treasury STRIP described above for \$93.42. We don't know what the price one year from today, $Price_1$, will be because we don't know how interest rates will change between now and then. So, we have to use our best guess of that price, which is its **statistical expectation** denoted $\mathbb{E}(Price_1)$. This expression reads "the expected value of $Price_1$ (i.e., the price one year from today)." The **expected return** of our zero coupon bond investment over the next year is

$$\mathbb{E}(r_{0,1}) = \frac{\mathbb{E}(Price_1)}{Price_0} - 1.$$

Because the expectation notation, \mathbb{E} , can be burdensome and it's usually clear whether we are talking about what *has* happened versus what *will* happen, we'll suppress it unless needed for clarity.

We'll talk more later about how to estimate expected values - i.e., put a number to them. For now we just need to understand that the expected value is a way for us to explicitly recognize that we don't know the value of future cash flows or discount rates, so we have to make an educated guess to estimate those values.

Of course, this is nothing new. Previous discussions about whether go to college and how to save for retirement (Chapter 2) and corporate decision making (Chapters 5 and 6) all relied on expected values. We don't know with certainty what our salary will be after we graduate college, just like we don't know for sure what the future free cash flows to a tablet project will be. Related, we don't know the future opportunity costs of those cash flows. We have to estimate all of these quantities, and we do so using expected values - our best guess given the information available to us. All we've done here is explicitly recognize that we have to estimate unknown future quantities so that we can distinguish between our guesses of future quantities (expectations) and what those future quantities eventually turn out to be (realizations).

7.3.8 Market Efficiency and Market Timing

If interest rates are high today relative to historical standards, is it a good time to buy bonds (i.e., lend) because they will likely come down in the future? Alternatively, if interest rates are low today relative to historical standards, is it a good time borrow (i.e., sell bonds)?

Market efficiency suggests the answers are no.

Market efficiency, loosely speaking, states that all available information relevant to the price of a bond (or any other asset traded in financial markets) is already incorporated in the price.¹³ Financial markets are extremely competitive. Many well-informed investors are trading financial assets to express their opinion of what the correct price should be. As a result, it's difficult to "beat the market." That is, it's difficult to make money above and beyond what's justified by the risk of the investment, because the market knows everything we do (and more) and has already incorporated that knowledge into the current price. Consider a couple of anecdotes to illustrate market efficiency in the context of interest rates.

In 2003, the Motley Fool, a financial advice website, noted "With interest rates still very low and housing prices continuing to appreciate steadily, there's a solid incentive for many people to no longer put off buying a house." One interpretation of this quote is that interested home buyers should buy now before interest rates go up in the future. But, if interest rates are going to go up in the future, wouldn't we think that finance professionals - bond traders and banks in particular - would know this before we would? It's unlikely we have information that professional fixed income investors don't.

What this means for any decision we make is that any expectation of future interest rate increases is already incorporated in the price of financial instruments. The seemingly low interest rate on our mortgage in 2003 probably isn't any lower than it's supposed to be given everything that the market knows. So, we're probably not getting a special deal on mortgages today just because the interest rate appears low relative to historical standards. To emphasize this point, figure 7.16 shows, interest rates were mostly flat for the next five years before declining for the next 13. As of 2021, many homeowners had mortgage rates below 3% on 30-year fixed-rate mortgages and below 2% on some adjustable-rate mortgages. In hindsight, 2003 was far from the best year for low interest rate loans.

Now consider what happened in December of 2008 after the onset of the Great Financial Crisis. The yield on the 90-day T-bill fell to zero. Certainly, that was a good time to bet that interest rates would increase and bond prices would fall. In fact, as figure 7.17 shows, the T-bill yield stayed at or near zero for the next seven years! Bond prices did not fall. The opposite, long-term bond prices rose as long-term interest rates decreased.

When we look at market interest rates and think, "these rates sure seem low (or high).

¹³Technically, there are three version of market efficiency differentiated by the ease with each investors can make profits from trading securities. Strong-form efficiency hypothesizes that investors cannot make profits trading on any information - public or private. Semi-strong form efficiency posits that investors can make profits trading on private information but not on on public information. Weak-form efficiency posits that investors can make profits trading on public and private information but not on information contained in past security prices.



Figure 7.16: 30-Year Fixed Rate Mortgage Average in the United States, 2003-2021 (Source: St. Louis Federal Reserve - Freddie Mac)

It’s a really good time to sell (or buy) bonds,” we need to pause and remember that the price of those bonds has already incorporated all of the information from professional investors. So, the real question to ask is: “what do we know that the professionals don’t?” Often, the answer is nothing. So, making investment decisions to take advantage of seemingly cheap or expensive financial instruments is often a fool’s game (no pun intended), unless we have a compelling story for why we believe the instrument is cheap or expensive *and* when the market will figure this out.

Let’s not confuse this discussion with the wisdom of shopping for a mortgage that we discussed in chapter 3. If we have decided to finance a home purchase, we want to make sure we pick the best mortgage. We don’t need to know where future interest rates are headed for this decision, only that a 4% 30-year fixed mortgage is a lot better than a 5% 30-year fixed mortgage, all other things equal. So, don’t confuse market efficiency as implying that shopping for a loan, whether it’s a home loan, an auto loan, a credit card, etc., doesn’t make sense. It does!

Market efficiency tells us to exercise skepticism when people say things like, “interest rates are low so now is a good time to buy a house or sell bonds,” or “the stock market is undervalued so we should buy stocks.” Market efficiency requires us to ask: What does this person know that the market doesn’t because most all the expectations about what may or may not happen in the future are already represented in the current price, which makes it difficult to earn profits beyond what we deserve for the risk we take.



Figure 7.17: Market Yield on U.S. Treasury Securities at 3-Month Constant Maturity, 2004-2018 (Source: St. Louis Federal Reserve - Board of Governors of the Federal Reserve System (US))

7.4 Investing in T-Notes and T-Bonds

7.4.1 Valuation

Valuing a coupon bond is another application of the Fundamental Value Relation. However, in light of our discussion of the term structure and yield curve, we need to recognize that each cash flow of a coupon bond may have a different discount rate. In other words, a more precise formulation of our Fundamental Value Relation acknowledges that r can vary with the time period.

$$Value_t = \frac{CashFlow_{t+1}}{1 + r_{t+1}} + \frac{CashFlow_{t+2}}{(1 + r_{t+2})^2} + \frac{CashFlow_{t+3}}{(1 + r_{t+3})^3} + \dots$$

Consider a two-year, \$100 par value T-note as of June 1, 2023 that matures on June 1, 2025 and has a **coupon rate** of 5%. This bond is referred to as the “5.0s of June 1, 2025,” denoting the coupon rate and maturity date. Like a zero coupon bond, the par value is paid at maturity on June 1, 2025. The periodic interest payments are determined by the coupon rate, 5%, which is expressed as an APR. Every year we get a total of $0.05 \times 100 = \$5$ of interest. However, because Treasury notes (and Treasury bonds) pay interest semi-annually, we get \$2.50 every every six months. The timeline of cash flows is shown in figure 7.18, which breaks out the coupon and principal payments.

To value this bond, we need to get the appropriate spot interest rates.

$$Value_0 = \frac{2.50}{(1 + r_1)} + \frac{2.50}{(1 + r_2)^2} + \frac{2.50}{(1 + r_3)^3} + \frac{102.50}{(1 + r_4)^4}$$

Date	6/1/23	12/1/23	6/1/24	12/1/24	6/1/25
6-month Periods	0	1	2	3	4
Coupons (\$)		2.50	2.50	2.50	2.50
Face (\$)					100.00
Cash Flows		2.50	2.50	2.50	102.50

Figure 7.18: Two-Year Treasury Note Cash Flows as of June 1, 2023

That is, we need the discount rates for 6-month (r_1), 12-month (r_2), 18-month (r_3), and ($24 - month$) risk-free loans. But, from where do we get these discount rates?

A hint comes from recognizing that a coupon bond is simply a collection - or **portfolio** - of zero coupon bonds with different face values. Our two-year T-note is equivalent to three zero coupon bonds with a face value of \$2.50 maturing in 6-, 12-, and 18-months, and a 24-month zero coupon bond with a face value of \$102.50. This equivalence is illustrated in table 5.

Security	Periods (Six Months)				
	0 ($Price_0$)	1	2	3	4
T-note	?	2.50	2.50	2.50	102.50
Replicating Portfolio					
6-month zero	?	2.50			
12-month zero	?		2.50		
18-month zero	?			2.50	
24-month zero	?				102.50
Portfolio	?	2.50	2.50	2.50	102.50

Table 5: Treasury Note and Replicating Portfolio

The cash flows of the portfolio of zero coupon bonds exactly match the timing and magnitude of the cash flows of the T-note. The investor receives \$2.50 every six months for one and a half years plus \$102.50 exactly two years from today. Because of this match, the portfolio of zero coupon bonds is referred to as a **replicating portfolio** for the T-note; the portfolio replicates the cash flows of another financial instrument.

We know how to value zero coupon bonds. We simply discount the future cash flow by the appropriate spot rate. There are several ways to estimate these rates, but perhaps the simplest, if not the most accurate, is to use the yields on Treasury bills and Treasury STRIPs, which correspond to spot rates. The yields for 6- and 12-month T-Bills and a two-

year STRIP from June 1, 2023 are presented in table 6. The 6-, 12-, and 24-month yields are from Treasury securities (Bills and STRIP) with corresponding maturities. The 18-month yield is estimated by linearly interpolating the 12- and 24-month yields.

Term (Months)	Yield (%)	Security
6	5.44	T-bill
12	5.11	T-bill
18	4.72	Linear interpolation of 12-month T-bill and 2-year STRIP
24	4.33	STRIP

Table 6: Treasury Bill and STRIP Yields as of June 1, 2023

With these yields we can value the T-bond using our Fundamental Value Relation.

$$\begin{aligned} Value_0 &= \frac{2.50}{1 + 0.0544/2} + \frac{2.50}{(1 + 0.0511/2)^2} + \frac{2.50}{(1 + 0.0472/2)^3} + \frac{102.50}{(1 + 0.0433/2)^4} \\ &= \$101.23 \end{aligned}$$

Each term in the expression above is the price of a zero coupon bond in our replicating portfolio. So, we can also value the T-note by valuing each zero coupon bond in the replicating portfolio and summing.

Security	Periods (Six Months)				
	0 ($Price_0$)	1	2	3	4
T-note	101.23	2.50	2.50	2.50	102.50
Replicating Portfolio					
6-month zero	$\frac{2.50}{(1+0.0544/2)} = 2.43$	2.50			
12-month zero	$\frac{2.50}{(1+0.0511/2)^2} = 2.38$		2.50		
18-month zero	$\frac{2.50}{(1+0.0472/2)^3} = 2.33$			2.50	
24-month zero	$\frac{102.50}{(1+0.0433/2)^4} = 94.08$				102.50
Portfolio	101.23	2.50	2.50	2.50	102.50

Table 7: Replicating Portfolio Valuation of Treasury Note as of June 1, 2023

What if the first coupon payment is not exactly 6-months from today? How does this affect our valuation of the bond? For example, imagine today is September 1, 2023 and our two-year T-note still matures on June 1, 2025. The first coupon arrives on December 1, 2023 - three months (91 days) from the date on which we are valuing the bond. The timeline is presented in figure 7.19.

Date	9/1/23	12/1/23	6/1/24	12/1/24	6/1/25
6-month Periods	0	0.5	1.5	2.5	3.5
Coupons (\$)		2.50	2.50	2.50	2.50
Face (\$)					100.00
Cash Flows		2.50	2.50	2.50	102.50

Figure 7.19: Two-Year Treasury Note Cash Flows as of September 1, 2023

We need to make two changes to our previous valuation. First, we need discount rates as of September 1, 2023, not June 1, 2023. We don't want to use stale data because interest rates may have changed since June 1, 2023. Second, we need discount rates for loan terms of 3-, 9-, 15-, and 21-months, or 0.5, 1.5, 2.5, and 3.5 six-month periods. In practice, we should be more precise in measuring the loan terms by using the number of days between cash flows - 91, 274, 457, and 639 - which we can convert to periods by dividing by the number of days in a six month period (e.g., 183). For the purpose of this example, let's ignore the rounding error.

Table 8 presents T-bill and STRIP yields as of September 1, 2023 for the new terms. Because the time to each cash flow other than the first do not coincide with terms of any zero coupon Treasury securities, we have to interpolate the yields for all but the three-month term.

Term (Months)	Yield (%)	Security
3	5.53	T-bill
9	5.42	Linear interpolation of 6- and 12-month T-bills
15	5.24	Linear interpolation of 12-month T-bill and 18-month STRIP
21	4.99	Linear interpolation of 18-month STRIP and 2-year STRIP

Table 8: Treasury Bill and STRIP Yields as of September 1, 2023

With the discount rates from Table 8 we can value the two-year T-note as of September 1, 2023 using our Fundamental Value Relation.

$$\begin{aligned}
 Value_0 &= \frac{2.50}{(1 + 0.0553/2)^{0.5}} + \frac{2.50}{(1 + 0.0542/2)^{1.5}} + \frac{2.50}{(1 + 0.0524/2)^{2.5}} + \frac{102.50}{(1 + 0.0499/2)^{3.5}} \\
 &= \$101.24
 \end{aligned}$$

Or, we can value each zero coupon bond in the replicating portfolio and then sum.

While it's convention to measure time in periods - half years for Treasuries - and use periodic interest rates to discount the bond cash flows, we could also measure time in years

Security	Periods (Six Months)				
	0 ($Price_0$)	1	2	3	4
T-note	101.24	2.50	2.50	2.50	102.50
Replicating Portfolio					
6-month zero	$\frac{2.50}{(1+0.0553/2)} = 2.47$	2.50			
12-month zero	$\frac{2.50}{(1+0.0542/2)^2} = 2.40$		2.50		
18-month zero	$\frac{2.50}{(1+0.0524/2)^3} = 2.34$			2.50	
24-month zero	$\frac{102.50}{(1+0.0499/2)^4} = 94.03$				102.50
Portfolio	101.24	2.50	2.50	2.50	102.50

Table 9: Replicating Portfolio Valuation of Treasury Note as of September 1, 2023

and use the effective annual rates to discount the bond cash flows. Using the EARs implied by the yields in Table 8, we get the same bond value as before.¹⁴

$$\begin{aligned}
 Value_0 &= \frac{2.50}{(1 + 0.0561)^{0.25}} + \frac{2.50}{(1 + 0.0549)^{0.75}} + \frac{2.50}{(1 + 0.0531)^{1.25}} + \frac{102.50}{(1 + 0.0505)^{1.75}} \\
 &= \$101.24
 \end{aligned}$$

7.4.2 Computing the Yield-to-Maturity

We can use the bond price and promised cash flows to find a coupon bond's yield to maturity, which, as a reminder, is the one discount rate such that the sum of the present values of all cash flows discounted at that rate equals the price of the bond - i.e., the bond's internal rate of return. Unlike the finding the yield for a zero coupon bond, we typically need the help of a computer to solve for the yield of a coupon bond.

Consider our two-year T-note from above as of June 1, 2023. The yield to maturity is the y that solves the following equation.

$$101.2257 = \frac{2.50}{(1 + y/2)} + \frac{2.50}{(1 + y/2)^2} + \frac{2.50}{(1 + y/2)^3} + \frac{102.50}{(1 + y/2)^4}$$

(Remember, yields are expressed as APRs and Treasuries are compounded semi-annually.) Using Excel's IRR or Goal Seek functions shows that the *periodic* yield is 2.1767%. The *bond equivalent yield* is $2 \times 2.1767\% = 4.3534\%$. Repeating this exercise for coupon bonds

¹⁴For example, the EAR for the 3-month yield is $(1 + 0.0553/2)^2 - 1 = 0.0561$.

of different terms enables us to construct a yield curve from coupon bonds, just as we did with zero coupon bonds.

Figure 7.20 presents the Treasury par yield curve as of June 1, 2023 reported by the U.S. Department of the Treasury. This yield curve is constructed from the bid prices of the most recently auctioned Treasury securities. While there are some technical aspects to constructing this curve, its interpretation is similar to what we've been discussing.¹⁵ Notice that the two-year treasury yield of 4.33% is very close to - within 2 basis points - what we computed above for our two-year T-note.




Figure 7.20: U.S. Treasury Par Yield Curve (June 1, 2023)

Caution: Odd First Coupons

Excel's IRR function assumes that the cash flows are equally-spaced in time. If this is not true, we cannot use IRR to find the yield of the bond. Consider our two-year note as of September 1, 2023. We need to find the y that solves the following equation.

$$101.2257 = \frac{2.50}{(1 + y/2)^{0.5}} + \frac{2.50}{(1 + y/2)^{1.5}} + \frac{2.50}{(1 + y/2)^{2.5}} + \frac{102.50}{(1 + y/2)^{3.5}}$$

¹⁵In practice, a **theoretical spot rate curve for Treasuries** is needed to provide discount rates to price Treasury securities because the Treasury does not issue zero-coupon bonds with maturities greater than one year. Construction of the such a curve is accomplished through a process called **bootstrapping** and is discussed in most fixed income textbooks.

There are three months until we receive the first coupon, but six months between every coupon thereafter.

One solution is to compute the value of the bond using one discount rate and then rely on Excel's GoalSeek function to find the discount rate that sets the value of the bond equal to the bond price.¹⁶ For our two-year T-note, the yield as of September 1, 2023 is 5.0031%, up from 4.3534% on June 1, 2023. The increase in yield is a consequence of the increase in interest rates between the two dates, as shown in figure 7.21. The Treasury yield curve shifted up.




Figure 7.21: U.S. Treasury Par Yield Curves (June 1 and September 1, 2023)

Despite this increase in interest rates, the value of the two-year T-note still increased from \$101.2257 on June 1, 2023 to \$101.2370 on September 1, 2023. But, if prices and yields move in opposite directions, how can this be? The reason is that yields weren't the only thing that changed between June and September first of 2023. The time to maturity decreased by three months, meaning the bond investor will be receiving their cash flows sooner. Had the bond yield *not* changed, the value of the bond on September 1, 2023 would have been

$$\begin{aligned}
 Value_0 &= \frac{2.50}{(1 + 0.043534/2)^{0.5}} + \frac{2.50}{(1 + 0.043534/2)^{1.5}} + \frac{2.50}{(1 + 0.043534/2)^{2.5}} + \frac{102.50}{(1 + 0.043534/2)^{3.5}} \\
 &= \$102.3215.
 \end{aligned}$$

¹⁶Excel also offers a function, XIRR, that can handle unevenly spaced cash flows.

So, the rising interest rates between June and September did in fact hurt the value of this bond.

Prices, Yields, and Coupons

When the price of a coupon bond is less than the face value of the bond, we say that the bond is **priced at a discount to par**. However, unlike zero coupon bonds, the price of a coupon bond can be greater than or equal to the par value. When the price of a coupon bond is greater than the face value of the bond, we say that the bond is **priced at a premium to par**. Finally, when the price of a coupon bond is equal to the face value of the bond, we say that the bond is **priced at par**.

Whether a coupon bond is priced at a premium, discount, or at par can be determined by the relation between the coupon rate and the yield. Bonds for which the coupon rate is greater than the yield are priced at a premium. Bonds for which the coupon rate is less than the yield are priced at a discount. And, bonds for which the coupon rate is equal to the yield are priced at par. (See the technical appendix for a proof of these results.)

Price Quotes

Table 10 presents price quotes for three of the over 300 Treasury notes and bonds quoted on J.P. Morgan's trading platform on September 1, 2023 . Buying and selling Treasury notes and bonds is no different from buying and selling T-bills and STRIPS, except the price quotes now recognize the coupon rate of the security. Otherwise, the information is the same.

Security Type	Coupon Rate (%)	Security Term	Maturity Date	Bid/Ask Price (per \$100 par)	Bid/Ask YTM (%)	Bid/Ask Quantity (\$ par)
Note	0.125	3-year	10/15/2023	99.36/99.48	6.06/4.90	25,000/40,000
Note	4.125	2-year	1/31/2025	98.40/98.77	5.35/5.04	25,000/40,000
Bond	3.625	30-year	2/15/2053	88.66/88.86	4.30/4.29	25,000/25,000

Table 10: Treasury Note and Bond Price Quotes as of September 1, 2023

For example, we can purchase up to 40,000 units, or \$40 million dollars of par value, of the 0.125s of 10/15/2023 at \$99.48 per \$100 of par value. These notes were issued almost three years ago on 10/15/2020. They will mature in 44 days, at which time we will receive the last coupon ($0.125\%/2 \times 100 = \$0.0625$) and principal (\$100). Because there is only one

more payment to be received, the yield to maturity can be computed using equation 7.1. For example, the ask yield is computed like so.

$$99.48 = \frac{100.0625}{(1 + y/2)^{44/183}} \implies y = \left(\left(\frac{100.0625}{99.48} \right)^{183/44} - 1 \right) \times 2 = 4.90\%$$

For the 30-year bond, the bid yield to maturity is computed by solving the following equation for y using Excel's GoalSeek.

$$88.86 = \frac{1.8125}{(1 + y/2)^{\frac{167}{184}}} + \frac{1.8125}{(1 + y/2)^{\frac{167}{184}+1}} + \frac{1.8125}{(1 + y/2)^{\frac{167}{184}+2}} + \dots + \frac{101.8125}{(1 + y/2)^{\frac{167}{184}+58}}$$

The solution is 4.32%, very close to what is reported in Table 10.¹⁷

7.4.3 Price Dynamics

Figure 7.22 presents the yield and price dynamics for the 4.0s of February 15, 2014, a ten-year Treasury note with a 4.0% coupon rate issued on February 15, 2004. The bond yield at issuance was 4.0479%, slightly higher than the coupon rate, which is why the bond is priced at a slight discount to par, \$99.8094. Let's discuss the similarities and differences of this figure with Figure 7.13 that presented the interest rate and price dynamics for zero coupon bonds.

First, more than one interest rate affects the price of a coupon bond in contrast to a zero coupon bond for which there is only one discount rate. To price our 10-year T-note at origination, we needed to use 20 possibly different interest rates corresponding to the timing of each coupon and the principal payment (i.e., 6-month, 1-year, 18-month, ..., 10-year). Rather than plotting all of these rates, the top part of figure 7.22 presents the bond yield, which is derived from the price and which captures the general level of the term structure at that time.

Second, when interest rates don't change (i.e., the yield on the bond is constant), the price of a coupon bond exhibits a saw-tooth pattern with no slope over time - the red series in the bottom graph. In contrast, a zero coupon bond price is an upward sloping line when there are no changes in interest rates. The saw-tooth pattern is a result of the coupons. As we approach a coupon payment, the price of the bond increases because we are getting

¹⁷The difference could be due to rounding or slight differences in the measurement of time.




Figure 7.22: Yield and Price Dynamics for a 10-Year Treasury Note

closer to receiving money. For example, the price of the bond the moment before receiving our first coupon is

$$2.00 + \frac{2.00}{(1 + 0.0479)} + \frac{2.00}{(1 + 0.0479)^2} + \dots + \frac{102.00}{(1 + 0.0479)^{19}} = \$101.6254,$$

which is \$2.02 greater than the price at which we purchased the bond, \$96.6094. The price is slightly larger than the first coupon because we are also closer to receiving all of the subsequent cash flows. Immediately after receiving the coupon, the price drops by the amount of the coupon. Remove the leading, undiscounted \$2.00 from the previous equation to reflect payment of the coupon. The price will fall from \$101.63 to \$99.63.

Third, the bond yield and price are negatively related like that for zero coupon bonds. Lower prices mean higher yields and vice versa. This relationship is mechanical in the

sense that it follows immediately from the Fundamental Value Relation; lower discount rates mean higher values. However, underlying the price changes are changes to the yield curve, i.e., interest rates. Figure 7.23 presents the treasury yield curve as of February 13, 2004, just before the bond was issued, and January 15, 2009, in the immediate aftermath of the financial crisis. Interest rates decreased significantly over this period and, as a result, bond prices increased. This change can be seen in figure 7.22 which indicates on January 15, 2009 a bond price of \$115.79 and bond yield of 1.13%. Had interest rates *not* changed, the bond price would have been \$101.71.




Figure 7.23: Treasury Yield Curves

Finally, like zero coupon bonds, coupon bond prices must eventually converge to their final payment at maturity, which equals the face value plus the last coupon. With relatively stable interest rates, this convergence mimics the sawtooth pattern.

7.4.4 Holding Period Returns

Computing the holding period return to a coupon bond is more challenging than doing so for a zero coupon bond because of the intermediate cash flows, i.e., the coupons. When we receive coupons, we have to reinvest them. But, because interest rates change over time, the rate of return earned on the reinvested coupons may differ from what is implied by the bond yield when we purchased it. Let's illustrate this using a two-year Treasury note with a 5% coupon that is currently priced at par.

The timeline for the bond is shown in figure 7.24. The periodic - six-month - bond yield, y , can be determined from the following equation.

$$100 = \frac{2.50}{(1+y)} + \frac{2.50}{(1+y)^2} + \frac{2.50}{(1+y)^3} + \frac{102.50}{(1+y)^4} \implies y = 2.50\%$$

We earn 2.50% per period, which implies a $2 \times 2.50 = 5\%$ bond yield. (We also could have recognized that because the bond is priced at par the yield equals the coupon rate.) Assuming interest rates don't change over the two years, our annual holding period return is $(1 + 0.0250)^2 - 1 = 5.0625\%$. And, over the two year life of the bond, our return would be $(1 + 0.050625)^2 - 1 = 10.38\%$.

6-month Periods	0	1	2	3	4
Coupons (\$)		2.50	2.50	2.50	2.50
Face (\$)					100.00
Cash Flows		2.50	2.50	2.50	102.50

Figure 7.24: Two-Year Treasury Note Cash Flows

Reinvestment Risk

But, as we've said before, the yield, which is just the internal rate of return for the bond, assumes that all intermediate cash flows are reinvested at the same rate of return as the internal rate of return. This means every time we receive a coupon, we have to find an investment that offers the same return as the bond did when we purchased it. Using our example, every time we receive a coupon payment of \$2.50, we must reinvest that money in something earning 5.0625% per year until the maturity of the bond. Of course, this is unrealistic because interest rates change frequently. Consequently, coupon bonds face **reinvestment risk**, in addition to interest rate risk. Not only do we have to worry about what happens to the price of our bond as interest rates change - should we decide to sell it before maturity - we have to worry about the interest rate at which we can reinvest the coupons.

Imagine that shortly after we buy the two-year T-note, its yield falls from 5% to 4% and stays there for the remaining life of the bond. Let's see what our investment return would be when we can only earn a 4% APR on the investment of our coupons? To answer this question, we can't use the internal rate of return or bond yield because the yield has changed. Instead, we have to use the following recipe to compute the **total return** of the bond.

1. Compute the future value of every coupon as of the maturity date, or date at which we intend to sell the bond, using the appropriate reinvestment rate(s).
2. Add the result from 1. to the principal value of the bond, or price at which we expect to sell the bond. This is the future value of all the money we expect to receive.
3. Take the ratio of the value computed in 2. to the current price of the bond and subtract one. Doing so will give us the holding period return of the bond, from which we can compute an effective annual or periodic return.

For our two-year T-note example,

1. The future value of the coupons two years from today invested at a 4% APR is

$$Value_4 = \frac{2.50}{0.04/2} \times ((1 + 0.04/2)^4 - 1) = \$10.3040$$

2. Adding the result from 1. to the principal value at maturity, \$100, equals \$110.3040.
3. Take the ratio of the result in 2. to the price of the bond and subtract one.

$$r_{0,4} = \frac{110.3040}{100} - 1 = 10.30\%$$

This is the holding period return over two years. The corresponding annualized return is $(1 + 0.1030)^{1/2} - 1 = 5.05\%$. This return is less than what is implied by the internal rate of return of the bond because the coupons were reinvested at a lower rate - 4% instead of 5%.

Figure 7.25 presents two plots corresponding to the 10-year T-note with a 4.0% coupon rate discussed above. The top plot presents the daily yield-to-maturity of the bond. The bottom plot presents two series. The red line represents daily returns. The blue line represents holding period returns, which are constructed from the daily returns like so.

$$r_{0,t} = (1 + r_{0,1}) \times (1 + r_{1,2}) \times \dots \times (1 + r_{t-1,t}) - 1$$

For example, the four day holding period return for the bond is equal to

$$r_{0,3} = (1 + 0.0008) \times (1 - 0.0007) \times (1 + 0.0014) - 1 = 0.15\%.$$

The figure shows that daily returns can be quite volatile, though nothing like stocks as we'll see in the next chapter.

Let's consider the returns to two different strategies with our 10-year, 4% coupon T-note to further illustrate the effect reinvestment risk on our returns.

1. Purchase the bond on February 15, 2004 and holds it until maturity 10 years later. The periodic internal rate of return on this strategy is just one half the yield to maturity, or approximately $4.05/2=2.025\%$. Compounding this six-month return 20 times over the life of the bond produces a ten-year holding period return of $(1 + 0.02025)^{20} - 1 = 49.32\%$, or $(1 + 0.4932)^{1/10} - 1 = 4.09\%$ per year. If instead we were forced to reinvest the coupons at the interest rate prevailing at the time of their distribution, our 10-year holding period return would have been 47.52%, or $(1 + 0.4752)^{1/10} - 1 = 3.97\%$ per year. Why is it lower than the IRR? Because as the top of figure 7.25 shows, interest rates declined over the investment period. So the coupons were reinvested at lower and lower interest rates.




Figure 7.25: Daily Bond Yields, Returns, and Holding Period Returns for a 10-Year 4.0% Treasury Note

2. Purchase the bond on June 13, 2006 for 95.23, and sells the bond on January 15, 2009

for 115.79. The timeline for this strategy is in figure 7.26. The bond equivalent yield of this strategy is 11.66%, which implies an annualized return of $(1+0.1166/2)^2 = 12.00\%$ that is substantially higher than the first strategy. Why? Simply put, we bought the bond at a low price when interest rates were high and sold the bond at a high price when interest rates were low. The capital gain from the sale far outweighs any lost interest on our coupons from having to reinvest them at lower rates.

Date	6/13/06	8/15/06	2/15/07	8/15/07	2/15/08	8/15/08	1/15/2009
6-month Periods	0	0.34	1.34	2.34	3.34	4.34	5.16
Coupons (\$)		2.00	2.00	2.00	2.00	2.00	
Sale Price(\$)							115.79
Cash Flows (\$)	-95.23	2.00	2.00	2.00	2.00	2.00	115.79

Figure 7.26: Cash Flows of Second Trading Strategy for 4.0s of February 15, 2014

There is a tradeoff when investing in coupon bonds. Falling interest rates increase the price of our bond but at the cost of a lower reinvestment rate. Which effect is more important depends on variety of factors including the coupon rate, the time to maturity, and the precise nature of interest rate changes.

7.5 Bond Arbitrage and Short Sales

Above we showed that a coupon bond is equivalent to a portfolio of zero coupon bonds. The portfolio of zeros whose cash flows exactly match the magnitude and timing of the coupon bond's cash flows was called a replicating portfolio. Because the cash flows between the replicating portfolio and the coupon bond are identical, their values should be equal. But, what happens when this isn't true?

An **arbitrage opportunity** exists. Arbitrage means risk-free profit. The beauty of an arbitrage opportunity is that, in theory, it doesn't require any money but still produces unlimited profit. By executing a set of transactions today - again with no money on our part - we receive money today. The best analogy I can think of for an arbitrage opportunity is finding money on the street with no one around. Let's illustrate this concept using the two-year T-note as of June 1, 2023 - the 5.0s of June 1, 2025 - that we valued at \$101.23.

Imagine that the bond was trading in the market for \$99.65, in which case the it's **mispriced**. Specifically, the bond is cheap relative to the price implied by the replicating portfolio of zeros, \$101.23. Why would this happen? Perhaps an institutional investor had to exit a large long position in the security and was forced to sell at a discount. What do we do when something is cheap? We buy it! Problem: We don't have any money. However, the replicating portfolio is expensive relative to the T-note. What do we do when something is expensive? We sell it! Problem: We don't own it. That is, we don't have any of the zero coupon bonds that make-up the replicating portfolio.

Here's where a little finance voo-doo comes in. We can borrow the zero coupon bonds from someone that does own them. Voila! Now we have zero coupon bonds that we can sell! When we do sell them, we receive \$101.12, the value of the replicating portfolio. Borrowing an asset and then selling it is called a **short sale**. Yes, it's legal, and yes, it happens quite frequently, though the devil is in the details we discuss below. With \$101.23 in our pocket, we can now purchase the cheap T-note for \$99.65, leaving us with the difference $101.23 - 99.65 = \$1.58$.

What happens next is we sit back and let the future cash flows cancel one another. Six months from today, the person from whom we borrowed the six-month zero is expecting to receive \$2.50. No problem because we will receive \$2.50 for the first coupon of the T-note we bought. We simply hand this money over to our security lender. Likewise, one year from today, the person from whom we borrowed the 12-month zero is expecting to receive \$2.50. We hand them the money from the second coupon payment of our T-note. This process continues until all of the cash flows we receive from the T-note are handed over to the people from whom we borrowed zero coupon bonds. The end result is that we made \$1.58 today with no money and never paid a penny in the future. This is an arbitrage whose strategy is detailed in Table 11.

Arbitrage opportunities are about **relative mispricing** where an asset is mispriced relative to one or more other assets. In our example, we don't know which asset is mispriced, the T-note or one or more of the zeros. We just know that they are mispriced relative to one another, and that's enough for us to make risk-free profit with no money. It's like finding a \$1.58 on the sidewalk and not having to worry about someone asking for the money back.

If we're smart, we'll scale this strategy by taking larger positions. Rather than short selling \$101.23 of the replicating portfolio, let's short \$10,123,000 of the replicating portfolio and use the proceeds to buy \$9,965,000 of the T-note. This leaves us with an arbitrage profit of \$157,572 today. But, this example begs the question: Why not scale this strategy up infinitely? We don't need any money!

Security	Periods (Six Months)				
	0 ($Price_0$)	1	2	3	4
Buy T-note	-99.65	2.50	2.50	2.50	102.50
Replicating Portfolio					
Short sell 6-month zero	$\frac{2.50}{(1+0.0553/2)} = 2.47$	-2.50			
Short sell 12-month zero	$\frac{2.50}{(1+0.0542/2)^2} = 2.40$		-2.50		
Short sell 18-month zero	$\frac{2.50}{(1+0.0524/2)^3} = 2.34$			-2.50	
Short sell 24-month zero	$\frac{102.50}{(1+0.0499/2)^4} = 94.03$				-102.50
Portfolio	101.2370	-2.50	-2.50	-2.50	-102.50
Difference	1.58	0	0	0	0

Table 11: Arbitrage Strategy of 5.0s of June 1, 2025

The problem comes when we consider what happens to the prices of the assets as we start executing our arbitrage strategy. By short selling the replicating portfolio, we are selling zero coupon bonds in the market. But, the more bonds we sell, the more the price of those bonds will decline. Likewise, we are buying 2-year T-notes. As we buy more, we will push up the price for these notes. In other words, as we scale up our arbitrage strategy, we will eat into our own arbitrage profits. In fact, it is because there are **arbitrageurs**, investors looking for arbitrage opportunities, that these opportunities are rare and short-lived. Taking advantage of an arbitrage opportunity is precisely what eliminates the arbitrage opportunity.

Short Selling

There are important challenges and limitations when short-selling securities.

1. Most large financial institutions have securities lending desks whose sole purpose is to lend out the securities they and their clients own. Nonetheless, we have to find someone willing to lend the security to us, which is not always possible.
2. Assuming we do find a willing lender, they're going to charge us interest on the loan. The more demand for the asset, the higher the interest rate on the loan. In fact, it may be that the interest expense on the security loan is high enough to make the arbitrage strategy unprofitable.
3. Lenders require **collateral**. When we borrow securities, we have to set aside cash or other securities in an account held by a third party to ensure the lender gets repaid.

in case the borrower (us) can't pay the lender back later. Typically, we have to post an amount equal to or greater than the amount we are short selling, depending on the quality of the collateral. Safer, less volatile assets like cash or Treasury's have relatively small *haircuts* - or discounts - than riskier, more volatile assets. If we borrow \$1 mil of securities, we may have to post \$1.03 million of Treasury securities or \$1.07 million of corporate bonds as collateral. The difference between the value of what we borrowed and the market value of the collateral we post is the haircut.

After we've borrowed and sold the assets, the worry isn't over.

1. We may receive a **margin call** demanding more collateral if (i) the value of the collateral we posted declines, or (ii) the price of the security that we borrowed increases. The purpose of margin calls is to ensure the lender is made whole should we default on our loan and not return the security. The lender can just keep the collateral, which is why they want to be sure they have enough collateral to cover any potential default by the borrower.
2. The lender may decide, within the terms of the lending agreement, to increase the interest rate and erode our arbitrage profits or even convert them to losses if demand for the security increases, i.e., more people want to borrow the security.
3. The lender may demand their security back before we are ready to return it. Perhaps they want to sell their security. In this case, we need to find someone else to lend us the same security to return to our original lender so we can maintain our arbitrage strategy. If we can't find another lender, we may get caught in a **short squeeze** in which we have to unwind our position, i.e., buy back the security in the market and return it to its owner. If this occurs after significant increases in the price of the security, this squeeze could be very costly to us.

In sum, short selling is risky in practice. This risk was on full display in January of 2021 when a number of hedge funds - institutional investors - had short sold GameStop stock in anticipation of sharp declines of its price. However, a large number of retail investors, led by Reddit user "Roaring Kitty," began buying GameStop stock in January of 2021. This buying had several implications beginning with a sharp increase in the price of GameStop stock. This price increase led to margin calls on investors with short positions (i.e., had short sold the stock). The buying frenzy also made it difficult to find stocks to borrow because owners were selling their shares. So, when short sellers tried to cover their short positions by

borrowing stock from another source, it was difficult. When they did find stock to borrow, the lending rate was high because of the demand from other short sellers. Consequently, many hedge funds lost a great deal of money from the price runup.

7.6 Interest Rate Risk

We've seen above that bond prices change over time in unpredictable ways because interest rates change over time in unpredictable ways. This unpredictability is referred to as **interest rate risk** or simply **rate risk**. Because changes in interest rates are the primary cause for changes in bond prices, understanding interest rate risk is critical for bond investors. Indeed, mismanaged or misunderstood interest rate risk that was largely responsible for several high profile financial failures including the bankruptcies of Orange County, California and Silicon Valley Bank.

7.6.1 Intuition

Consider figure 7.27 which shows how the prices of four different bonds change when interest rates change by different amounts. The horizontal axis measures the percentage change in interest rates (e.g., 5% to 6% is +1%). The vertical axis measures the percentage change in the price (e.g., \$100 to \$95 is -5%). The picture highlights how different features of a bond - maturity, coupon rate, and initial yield - affect the interest rate sensitivity of the bond, that is, how reactive bond prices are to changes in interest rates.

Several lessons can be gleaned from this figure.

1. **Bond prices and yields are inversely related.** We've seen this, but it's worth repeating. All of the curves slope down. When interest rates go up (down), bond prices go down (up).
2. **Bond prices become more sensitive to interest rate changes as the maturity of the bond increases.** The slope of curve A - the bond with the shortest maturity - is the flattest at every point.
3. **Bond prices become less sensitive to interest rate changes as the coupon rate increases.** The slope of curve B (zero coupon) is steeper than that for curve C (12% coupon).




Figure 7.27: Bond Price Sensitivity to Interest Rate Changes

4. **Bond prices are more sensitive to decreases in interest rates than otherwise similar increases in interest rates.** Each curve gets progressively less steep as we move from left to right (i.e., as interest rates increase). This property is referred to as **convexity**, but it only applies to bonds without embedded options.

It is important to point out that each of these properties are “all else equal,” meaning that we’re only changing one feature of a bond at a time. For example, bond prices are less sensitive to interest rate changes as the coupon rate increases assuming the maturity and yield are the same.

Now let’s see how to measure or quantify interest rate risk.

7.6.2 DV01

Let’s start by changing the interest rate a little and compute by how much the bond price changes. Consider a 7-year zero coupon bond with a \$100 face value and an annually compounded yield of 10%. The price of the bond today is

$$Price_0 = \frac{100}{(1 + 0.10)^7} = \$51.32.$$

Now let's compute the price if the interest rate decreases by 1 basis point (bp), or 0.01%.

$$Price_0^* = \frac{100}{(1 + 0.0999)^7} = \$51.35$$

A one basis point decrease in the interest rate generates a $51.35 - 51.32 = \$0.0327$ increase in the price. This value is called **dollar value of a one basis point decrease** or **DV01** for short. It measures interest rate risk by quantifying the change in value of our bond in response to a small change in the interest rate.

Of course, we can estimate the bond price change in response to any magnitude change in interest rates. For example, if interest rates decrease by one point, 1%, the bond price will be

$$Price_0^* = \frac{100}{(1 + 0.09)^7} = \$54.70.$$

The dollar value of a one percentage point decrease in the interest rate is $54.70 - 51.32 = \$3.38$.

This estimate is pretty close to $DV01 \times 100 = \$3.27$.¹⁸ We can approximate the change in bond price for any sized change in interest rate by multiplying $DV01$ by the interest rate change measured in basis points. For example, to know how the bond price will respond to a 0.5% decrease in interest rates, we can compute

$$DV01 \times \text{Interest rate change (bps)} = 0.0327 \times 50 = \$1.635,$$

where **bps** stands for basis points. *Warning:* The larger the interest rate change, the worse this approximation.

What would the DV01 for our bond be if its maturity were 20 years instead of seven?

$$\begin{aligned} Price_0 &= \frac{100}{(1 + 0.10)^{20}} = \$14.86436 \\ Price_0^* &= \frac{100}{(1 + 0.0999)^{20}} = \$14.89141 \\ \implies DV01 &= \$0.02705 \end{aligned}$$

DV01 is smaller with the longer maturity, which seems to contradict our intuition that longer term bonds tend to be more sensitive to interest rate changes. However, the vertical axis in figure 7.27 measures the *relative* or *percent* price change. Let's divide each DV01 measure by the respective bond price and multiply by 100 to change the interpretation from a one

¹⁸The difference is due to the nonlinear relation between the price and yield.

basis point to a one percentage point change in the yield.

$$\begin{aligned} \text{7-year bond: } & \frac{0.0327 \times 100}{51.32} = 0.0637 \\ \text{20-year bond: } & \frac{0.0271 \times 100}{14.86} = 0.1817 \end{aligned}$$

A one percentage point decrease in interest rates leads to a 6.37% increase in the price of the 7-year bond, but an 18.17% increase in the price of the 20-year bond. The longer-term bond price is almost three times more sensitive than that of the shorter-term bond.¹⁹

What about a coupon bond? To value a coupon bond we need to discount each cash flow by a possibly different interest rate that we obtain from the yield curve. How can we measure this bond's interest rate risk since it is subject to more than one interest rate? One way to approximate interest rate risk in this case is to assume that the term structure of interest rates is flat - i.e., there is only one interest rate for all terms - and equal to the current bond yield. In this case, we can answer the question: If the term structure is flat and it shifts up by a small amount, by how much does the price of the bond change?

Take the 5.0s of June 1, 2025 as an example. The price and yield of the bond as of June 1, 2023 are \$108.38 and 4.35%, respectively. To compute DV01 for this bond, we start by revaluing the bond at a yield that is one basis point lower than its current yield.

$$\begin{aligned} Price_0^* &= \frac{2.50}{(1 + 0.0434/2)} + \frac{2.50}{(1 + 0.0434/2)^2} + \frac{2.50}{(1 + 0.0434/2)^3} + \frac{102.50}{(1 + 0.0434/2)^4} \\ &= \$101.2448 \end{aligned}$$

The DV01 for our coupon bond is $101.2257 - 101.2448 = \$0.0191$. Relative to the price of the bond, a one percentage point decrease in interest rates leads to a $0.0191/101.2257 \times 100 = 1.8877\%$ increase in the bond price.

Let's recompute the DV01 for the same bond assuming it's coupon rate was 10% instead of 5%. (The different coupon rate will impact the price and yield, even though we're assuming the same term structure of interest rates.)

$$\begin{aligned} Price_0 &= \frac{5.00}{(1 + 0.0437/2)} + \frac{5.00}{(1 + 0.0437/2)^2} + \frac{5.00}{(1 + 0.0437/2)^3} + \frac{105.00}{(1 + 0.0437/2)^4} \\ &= \$110.6623 \\ Price_0^* &= \frac{5.00}{(1 + 0.0436/2)} + \frac{5.00}{(1 + 0.0436/2)^2} + \frac{5.00}{(1 + 0.0436/2)^3} + \frac{105.00}{(1 + 0.0436/2)^4} \\ &= \$110.6825 \\ \implies DV01 &= \$0.0202 \end{aligned}$$

¹⁹Notice that the maturity of the long-term bond is almost three times larger than that of the short-term bond - 20 versus 7. This relation between the maturity of the bond and its interest rate sensitivity is not a coincidence.

Relative to the price of the bond, a one percentage point decrease in interest rates leads to a 1.8294% increase in the bond price. Consistent with the intuition from above, the higher the coupon rate, the less sensitive the bond is to changes in interest rates, though the difference is small.

As a quick aside, the differences in our example may appear small. Who cares if we lose \$0.02, for example, when interest rates increase a little? Most bond issues are in the millions, if not billions, of dollars. The U.S. treasury issued \$15 trillion in 2022. So, small changes in interest rates can have economically large value effects.

7.6.3 Duration

Another measure of interest rate sensitivity is **Macaulay Duration** or **Duration**. Duration has a scary looking formula with an interesting interpretation that is closely related to the intuition we just discussed. Let's annotate the formula and write it two different ways to highlight its interpretations.

$$\begin{aligned}
 \text{Duration} &= \frac{1}{\text{Price}_0} \left(\frac{\overbrace{CF_1}^{PV(CF_1)}}{(1+y/k)^1} \cdot \frac{1}{k} + \frac{\overbrace{CF_2}^{PV(CF_2)}}{(1+y/k)^2} \cdot \frac{2}{k} + \dots + \frac{\overbrace{CF_T}^{PV(CF_T)}}{(1+y/k)^T} \cdot \frac{T}{k} \right) \\
 &= \underbrace{\frac{PV(CF_1)}{\text{Price}_0}}_{\text{weight}_1} \cdot \frac{1}{k} + \underbrace{\frac{PV(CF_2)}{\text{Price}_0}}_{\text{weight}_2} \cdot \frac{2}{k} + \dots + \underbrace{\frac{PV(CF_T)}{\text{Price}_0}}_{\text{weight}_T} \cdot \frac{T}{k} \tag{7.5}
 \end{aligned}$$

In the above equations, y is the bond yield to maturity expressed as an APR and k is the compounding frequency (times per year).

Duration tells us the weighted average of the waiting times for receiving a bond's cash flows. Duration units are years and it measures the *effective* maturity of the bond. Each weight is the ratio of the present value of the cash flow to the price. Because the price is the sum of all the present values, the weights sum to one. The terms $1/k, 2/k, \dots, T/k$ correspond to the time measured in years to each cash flow. An example will be helpful.

Table 12 details the duration computation for our 2-year T-note from earlier. The price and yield of the bond are repeated for ease of reference. The cash flows correspond to the coupon and principal payments. The present value of each cash flow is computed assuming semi-annual compounding to coincide with the frequency of the cash flows. For example, the present value of the third coupon payment is

$$\frac{2.50}{(1 + 0.04354/2)^3} = \$2.34$$

	Periods (Six-Months)				
	0	1	2	3	4
Price (\$)	101.23				
Yield (%)	4.3534				
Cash flows (\$)		2.50	2.50	2.50	102.50
Present values, $PV(CF_t)$		2.45	2.39	2.34	94.04
Weights, $PV(CF_t) / \text{Price}$		0.02	0.02	0.02	0.93
Time to cash flow (Years)		0.50	1.00	1.50	2.00
Weight x Time		0.01	0.02	0.03	1.86
Sum, Duration	1.93				

Table 12: Duration of Two-year T-note: 5.0s of June 1, 2025

The weights are computed by scaling each present value by the price of the bond. The time to cash flow row shows the time measured in years until each cash flow is received. For example, period 3 is 1.5 years from today.

Summing the products of the weights and the times (2nd to last row) gives us a duration estimate for the bond of 1.93 years. This number is a little less than the maturity of the bond because the coupon payments shift the distribution of cash flows forward in time. When there are no coupon payments, i.e., a zero coupon bond, duration equals maturity.

While interesting, duration is measured in years. So, it's not immediately clear how this helps us understand the interest rate sensitivity of our bond's price. To make this measure more meaningful, we focus on a variation of Macaulay duration called **modified duration**.

$$\text{Modified duration} = \frac{\text{Duration}}{1 + y/k} \quad (7.6)$$

For our 2-year T-note, the modified duration is

$$\frac{1.93}{1 + 0.04354/2} = 1.8874.$$

Modified duration tells us the percentage change in the price of the bond for a one percentage point (100 basis point) decrease in the bond yield.²⁰ More precisely, it measures the impact on the bond price of the entire yield curve shifting down by 1%. For our 2-year T-note, a one percentage point decrease in interest rates leads to a 1.89% increase in the value of the bond. As long as yields are relatively small, i.e., $(1 + y/k)$ is close to one,

²⁰Modified duration is an example of a **semi-elasticity**, which measures the percentage change of a function in response to a one unit change in a variable of that function.

duration and modified duration are quantitatively similar, leading practitioners to interpret Macaulay duration in a similar manner.

There is a close relation between DV01 and modified duration. DV01 is the dollar value of one basis point. Let's multiply this by 100 to get the dollar value of a point (i.e., 1%), $\$0.01911 \times 100 = \1.911 . If we divide this by the price of the bond, \$101.23, we get an estimate of the percent change in the price of the bond for a one percent change in the yield, $1.911/101.23 = 1.8877\%$ - similar to our modified duration estimate. Mathematically,

$$\frac{DV01 \times 100}{Price} \times 100 \approx -\text{Modified Duration}$$

See the technical appendix for further details.

7.7 Liability Driven Investment and Immunization

Let's revisit LDI by imagining our child was accepted into the University of Pennsylvania and has decided to defer enrollment for one year, meaning they will start school two years from today. Let's also assume that tuition, room, and board will be \$80,000 at the start of school and will grow by 5% per year for the next three years. How can we invest today to ensure we will have enough money to cover our child's tuition expenses?

Using a cash flow matching approach, we should buy four zero-coupon bonds with 2-, 3-, 4-, and 5-year maturities and whose face values equal our future expenses. This investment approach is displayed in table 13. Tuition expenses are in the second row, under which is the term structure of bond equivalent yields assuming semi-annual compounding. The prices of the zero coupon bonds are simply the present value of the par values. For example, the price of the 2-year zero is

$$\frac{80,000}{\left(1 + \frac{0.0494}{2}\right)^4} = \$72,560.96.$$

Buying a zero coupon bond with a par value of each future liability exactly replicates the magnitude and timing of the future liabilities. No matter what interest rates do, we will have enough money to make all of the future tuition payments as long as the bonds are repaid in full. Thus, the zero coupon bonds should be low risk, like Treasuries or very high quality corporate bonds, to match the risk of our tuition payments.

Year	0	1	2	3	4	5
Tuition			80,000.00	84,000.00	88,200.00	92,610.00
Yield (%)		5.40	4.94	4.66	4.52	4.38
2-year	-72,560.96		80,000.00			
3-year	-73,157.94			84,000.00		
4-year	-73,760.23				88,200.00	
5-year	-74,571.67					92,610.00
Portfolio	-294,050.81		80,000.00	84,000.00	88,200.00	92,610.00

Table 13: Liability Driven Investing with Zero Coupon Bonds - Saving for College

7.7.1 Duration Matching

Two features of our assets - the portfolio of zero coupon bonds - should be noted. First, our portfolio value is equal to the present value of the future liabilities, \$294,050.81. Second, the (modified) duration of our portfolio is equal to the duration of our liabilities - 6.69. If interest rates change by 100 basis points, the value of our liabilities will change by 6.69% or \$19,666.57. But, so will the value of our assets because they are worth the same amount and have the same duration. Because any change in the value of our liabilities due to interest rate changes is met with an equivalent change in our assets, we will always be able to fund our liabilities. In other words, our funding model is **immunized** against interest rate changes.

The two conditions for interest rate immunization are:

1. The present value of our assets equals the present value of liabilities, and
2. The duration of our assets equals the duration of our liabilities.

These are pretty general conditions and permit a lot of different asset portfolios other than our zero coupon bond strategy. For example, imagine we wanted to invest in a three-year and a 20-year coupon bond whose durations are 2.73 and 8.71. How much of each bond would we need to buy?

The duration of a portfolio is just the market value weighted average of the component durations. (Sorry for that mouthful.). In other words, if our portfolio has N bonds whose values are v_1, \dots, v_N and whose durations are D_1, \dots, D_N , then the portfolio duration, D_{port} , is as follows.

$$D_{port} = \frac{v_1}{v_1 + \dots + v_N} D_1 + \frac{v_2}{v_1 + \dots + v_N} D_2 + \dots + \frac{v_N}{v_1 + \dots + v_N} D_N \quad (7.7)$$

Note, the weights sum to one.

Equation 7.7 implies that the duration of our portfolio can be written like so.

$$w_A \times D_A + w_B \times D_B$$

We don't know the weights, w_A and w_B , but we know they must sum to one. We also know that the portfolio duration should equal the duration of our liabilities, 6.69. Putting this together means we can figure out the weights.

$$6.69 = w_A \times 2.73 + (1 - w_A) \times 8.71 \implies w_a = 33.82\%$$

One minus the weight on bond A is the weight on bond B, 66.18%. Because the value of the portfolio should equal the present value of our liabilities, \$294,050.81, we need to buy $0.3382 \times 294,050.81 = \$99,450.41$ of bond A and $0.6618 \times 294,050.81 = \$194,600.39$ of bond B.

There are some limitations with this hedging approach. Specifically, our asset portfolio will only be immunized against interest rate changes if (i) the yield curve is flat and (ii) any changes in the yield curve are parallel shifts. In other words, every interest rate is the same and all subsequent changes are by the same amount. Clearly, these conditions are unrealistic. So, there is risk associated with immunizing a portfolio against interest rate changes, a full discussion of how to deal with this immunization risk is deferred to more specialized texts.²¹

7.8 Default Risk

Default risk or **credit risk** refers to the possibility that a borrower will not pay their obligation. If a borrower does not make a promised payment, such as a coupon or principal payment, the borrower is said to be in **payment default** or just **default**.²² When there is a chance the borrower can't repay a loan, the future cash flows are uncertain or risky. We address this risk by using *expected* cash flows, just like we used expected cash flows in virtually every other application we've tackled: saving for retirement, value college, capital budgeting, etc. Put simply, we have to estimate future bond cash flows given the likelihood the borrower will default.

Consider a simple example. A one year zero coupon bond is currently priced at \$92 per \$100 of par value. Assume there is a 10% probability that the borrower will default and not

²¹ "Bond Markets, Analysis, and Strategies" by Frank J. Fabozzi and Francesco A. Fabozzi is an excellent reference.

²² **Technical default** occurs when a borrower violates a covenant other than one requiring the payment of interest and principal. For example, a borrower's leverage ratio may breach a certain level or a borrower may sell an asset or acquire a company they are not supposed to as stated in the bond contract.

repay the \$100 one year from today. Also assume that the **recovery rate** on the bond is 60%, meaning that if the borrower defaults we can expect to recover 60% of what is owed, or $0.60 \times 100 = \$60$. When borrowers default, lenders often recover some fraction of what they are owed by seizing the borrower's assets.

The yield to maturity of this bond is computed using the *promised* cash flow, \$100.

$$\begin{aligned} Price_0 = \frac{CF_1}{(1+y)} &\implies y = \frac{CF_1}{Price_0} - 1 \\ &= \frac{100}{92} - 1 = 0.087. \end{aligned}$$

That is, the yield on this bond is 8.7% per annum. However, because the bond is risky, the yield and the expected return - what investors actually expect to earn - are different.

The cash flow one period from today, CF_1 , is a **random variable**. We don't know what value it will take one year from today. We only know its **probability distribution**. That is, we know what values the random variable can take and the probability of each value occurring. With this information, we can compute the expected value of next year's cash flow in our valuation relation. We do this by multiplying each value the variable can take by its corresponding probability and summing. In our bond example, the expected cash flow one year from today is

$$\begin{aligned} \mathbb{E}(CF_1) &= Pr(\text{Default}) \times \text{Cash flow in default} + Pr(\text{No default}) \times \text{Cash flow no default} \\ &= 0.10 \times 60 + (1 - 0.10) \times 100 \\ &= \$96. \end{aligned}$$

Using the *expected* cash flow, as opposed to the promised cash flow, allows us to compute the expected return or discount rate for the bond.

$$r = \frac{\mathbb{E}(CF_1)}{Price_0} - 1 = \frac{96}{92} - 1 = 0.043$$

The expected return on our bond is 4.3%, significantly less than the 8.7% yield. The reason for this difference is that we - the investor - don't expect to receive the promised \$100 a year from today. We only expect to receive \$96 because (i) there is a chance the borrower will default, and (ii) if they do default we will not recover the entire amount owed. See the technical appendix of this chapter for a more detailed example.

There is an important lesson here for corporate managers as well as investors. The debt cost of capital for a risky company is *not* the yield on its debt. The yield overstates the expected return on debt, i.e., the debt cost of capital, because it ignores default risk. Unless

a company is financially sound, the yield to maturity and expected return can be quite different as our simple example shows.

Related, the expected return on a bond with default risk is *greater* than the risk-free rate. Bond investors demand a **risk premium**, or extra return above and beyond the risk-free rate to compensate them for the risk of their cash flows. Formally, we can think of the expected return on a bond - any risk asset actually - as comprised of a (1) risk-free or guaranteed return that compensates us for the time value of money, and (2) a risk premium that compensates us for any risk or uncertainty in the cash flows.

$$r = \text{Risk-free return} + \text{Risk premium} \quad (7.8)$$

In our defaultable bond example, the expected return of 4.3% is less than the bond yield, but it will be greater than the risk-free return, such as the return on a similar maturity Treasury, because of the risk premium investors demand.

Figure 7.28 presents the daily time series of **credit spreads** from 1996 to 2024. These figures represent the risk premium component from equation 7.8 and correspond to the additional interest - above that earned on similar maturity Treasuries - investors earned on a portfolio of relatively safe corporate bonds. The average credit spread over the sample period is 1.51%. However, the spread exhibits a great deal of volatility, especially around the great financial crisis of 2008-2009 and the Covid-19 pandemic in early 2020. During these periods, investors ran from risky investments - e.g., stocks, real estate, and corporate bonds - and invested in safe Treasuries. The result was a decline in corporate bond prices and sharp rise in credit spreads.

7.8.1 Credit Ratings

We saw above that computing interest rate risk is relatively easy, if not tedious. DV01 and duration provide quick and reasonably accurate estimates of the change in value of our bond or bond portfolio that occurs when interest rates change. Estimating default risk is more difficult. What matters are both the probability that a borrower defaults *and* the **loss given default**. That is, if the borrower defaults, how much money do we lose or alternatively what is our recovery rate? If we don't lose anything when a borrower defaults, then default presents no risk because we don't lose any money.

To estimate the probability and loss given default, we have to understand the financial risk of the issuing entity (e.g., company, government), as well as its incentives. This process is difficult and requires a great deal of time, energy, and expertise. In response, rating agencies




Figure 7.28: ICE BofA US Corporate Index Option-Adjusted Spread, Percent, Daily, Not Seasonally Adjusted (Source: St Louis FRED)

evolved to distill default risk into easy to interpret metrics called **credit ratings**. Companies like Moody's, Standard and Poors (S&P), and Fitch publish credit ratings and other information on tens of thousands of fixed income investments and hundreds of companies to help investors understand the default risk of their investments.

Rating scales are depicted in table 14, which is adapted from Wikipedia.²³ These are ratings for long-term bonds, where long-term means in excess of one year. Short-term bonds, such as commercial paper, have slightly different scales. Ratings fall into two broad groups: investment grade and below investment grade. The former, also known as high quality, consist of relatively safe bonds that are unlikely to default. The later, also known as speculative grade, high yield, and junk, are much more likely to default as indicated by their average default probabilities obtained from

These ratings and the default risk they measure manifest themselves in corporate credit spreads. Figure 7.29 presents credit spreads by different credit ratings. Spreads increase as we move down the ratings scale. Riskier borrowers are charged higher interest rates. If we look carefully, we can see a particularly large gap in credit spreads between BBB and BB

²³The ratings from Fitch overlap with those from S&P except that there is no "SD" in the Fitch rating scale. The default probabilities are weighted long-term average one-year global default probabilities from Standard & Poor's "Default, transition, and recovery: 2023 annual global corporate default and rating transition study."

Moody's	S&P and Fitch	Average Default Probability (%)	Description
Investment Grade			
Aaa	AAA	0.00	Prime
Aa1	AA+	0.02	High grade
Aa2	AA		
Aa3	AA-		
A1	A+	0.05	Upper medium grade
A2	A		
A3	A-		
Baa1	BBB+	0.14	Lower medium grade
Baa2	BBB		
Baa3	BBB-		
Below Investment Grade			
Ba1	BB+	0.57	Non-investment grade
Ba2	BB		
Ba3	BB-		
B1	B+	2.98	Highly speculative
B2	B		
B3	B-		
Caa1	CCC+	25.98	Substantial risk
Caa2	CCC		
Caa3	CCC-		
Ca	CC		Extremely speculative
	C		Default is imminent
C	RD		In default
	SD		
	D		

Table 14: Credit Ratings Scales and One-Year Weighted Long-Term Average Global Default Probabilities.

ratings, which are on different sides of the investment grade-speculative grade divide. Many institutional investors (e.g., banks, insurance companies, pension funds) have limitations on investments in high yield debt, which reduces demand for these bonds and therefore bond prices. The decline in price means higher yields all else equal.

While useful, credit ratings have their limitations. They are slow to incorporate new information relative to bond prices. They can also be inaccurate, as they were for many




Figure 7.29: ICE BofA US Corporate Index Option-Adjusted Spread, Percent, Daily, Not Seasonally Adjusted (Source: St Louis FRED)

securitizations that defaulted during the 2008 financial crisis.²⁴ That said, ratings of corporate bonds have been relatively consistent over the years and they remain a useful guide for assessing the default risk of many fixed income instruments.

7.9 Inflation Risk

In addition to interest rate and credit risk, fixed-rate bonds - bonds whose coupon rate is constant over the life of the bond - face **inflation risk**. Fixed coupon payments promise a nominal rate of return, not a real rate of return. For example, buying a treasury note with a coupon rate of 4% will generate \$2 of interest income per \$100 of par value every six months for the life of the note. If inflation is 8% per year, as it was in the US in September of 2022, the real rate of return is approximately $4 - 8 = -4\%$ per year. This result means that while we are receiving money from the bond, we are losing purchasing power. We can't buy as

²⁴In simple terms, a securitization starts with a bunch of loans, such as auto loans to consumers or commercial loans to companies. These loans are brought together using proceeds from investors to form a **pool of loans**. The investors who contribute the money to buy the loans and form the pool are given bonds whose payments come from the repayment of loans in the pool. So, as consumers pay their car loans or companies pay their loans, those payments are passed on to the bond investors. The auto loan example is referred to as an **asset-backed securitization** or **ABS**, the commercial loan example is a **collateralized loan obligation** or **CLO**.

many goods and services because the price is going up more quickly than we are earning interest.

7.9.1 TIPS

In 1997, **Treasury Inflation Protected Securities** or (**TIPS**) were introduced to address the shortcomings of **nominal bonds**, or bonds exposed to inflation risk. The coupon rates of TIPS are fixed and coupons are paid semi-annually, just like Treasury notes and bonds. However, the principal of TIPS are tied to the prices of goods and services in the economy.²⁵ As prices increase or decrease, so too does the principal of the TIPS. Because the principal is changing, the coupon payments will change even though the coupon rate is fixed. Let's see how this works with an example.

Date	CPI	Index Ratio	Unadjusted Principal	Real Cash Flow	Adjusted Principal	Nominal Cash Flow
10/15/21	273.2577	1.00000	1,000.00		1,000.00	
4/15/22	282.3464	1.03326	1,000.00	0.6250	1,033.26	0.6458
10/15/22	296.2286	1.08406	1,000.00	0.6250	1,084.06	0.6775
4/15/23	308.1380	1.12765	1,000.00	0.6250	1,127.65	0.7048
10/15/23	323.2955	1.18312	1,000.00	0.6250	1,183.12	0.7394
4/15/24	330.2265	1.20848	1,000.00	0.6250	1,208.48	0.7553
10/15/24	332.9618	1.21849	1,000.00	0.6250	1,218.49	0.7616
4/15/25	328.2344	1.20119	1,000.00	0.6250	1,201.19	0.7507
10/15/25	305.7289	1.11883	1,000.00	0.6250	1,118.83	0.6993
4/15/26	328.5623	1.20239	1,000.00	0.6250	1,202.39	0.7515
10/15/26	340.7086	1.24684	1,000.00	1000.6250	1,246.84	1247.6193
Real yield				(1.7785)		
Nominal yield				1.1300		
Implied inflation				2.9085		

Table 15: 5-year, 0.125% Coupon TIPS Example Cash Flows

Table 15 shows how a 5-year, 0.125% coupon TIPS works. The security was **dated** as of October 15, 2021, meaning the clock starts ticking down towards maturity as of that date. To adjust for inflation, the **consumer price index** **consumer price index (CPI)**, which

²⁵More precisely, the principal of the securities is indexed to the **Consumer Price Index for All Urban Consumers (CPI-U)**.

measures the general level of prices in the economy, is used. The CPI as of the dated date was 273.2577 and is referred to as the **reference CPI**.²⁶ The index ratio is the ratio of the current CPI to the reference CPI and is used for adjusting the principal of the bond.

The first coupon payment is on April 15, 2022. For a Treasury note, we saw above that the coupon is computed by multiplying the par value of the bond times one half the coupon rate - remember the coupons are paid semi-annually. In this example, a Treasury note with the same coupon rate as the TIPS would pay a coupon equal to $1,000 \times 0.00125/2 = \0.625 per \$1,000 of face value. This value is the *real* cash flow paid to investors, where real refers to deflated or inflation adjusted, not what investors actual receive. (Annoying terminology, I know.)

What TIPS investors actually receive is the *nominal* cash flow or coupon payment. This value is computed by first multiplying the principal of the TIPS by the index ratio to obtain an adjusted principal figure. On April 15, 2022, the adjusted principal of the bond was $1,000 \times 1.03326 = \$1,033.26$. The coupon is therefore $1,033.26 \times 0.000125/2 = \0.65 per \$1,000 of par value, slightly higher than that for an otherwise similar Treasury security.

Notice that by multiplying the principal by the index ratio, we are adjusting the principal of the bond, and by extension the coupon payment, by the rate of inflation. The price level went up from 273.2377 to 282.3464 between October 15, 2021 and April 15, 2022. This increase in prices is an inflation rate of $282.3464/273.2577 - 1 = 0.03326$, or 3.326%. To keep up with this price increase, we need more money. To ensure we are just as well off before the price increase, we need 3.326% more money, and that is exactly what the principal adjustment does. It increases the principal, and by extension the coupon payment, by 3.326%, the rate of inflation. So, while our nominal cash flows - what we're actually getting - are increasing, our real cash flows - what we can buy - are staying constant.

7.10 Term Structure of Interest Rates

We introduced the term structure of interest rates above in the context of the yield curve. As a reminder, the term structure is the relation between loan terms and interest rates while the yield curve is an illustration of the term structure. Because there are many different types of loans, there are many term structures and corresponding yield curves. Figure 7.30 presents several examples in three separate graphs.

²⁶A couple of details worth mentioning. The bond was issued on October 29, 2021. The CPI is computed using historical values and a linear interpolation scheme. See the Treasury Direct website for more details.




Figure 7.30: Yield Curves for Different Markets, Yield Curves for Different Credit Ratings, and Treasury Yield Curves at Different Points in Time. Sources: Treasury.gov, Tradeweb, Bloomberg, and NASDAQ

The top left graph presents three different yield curves corresponding to three different bond markets as of June 2022: Treasury, Municipal, and High Quality Market (HQM) Corporate. Each market has a distinct term structure differing in level and slope, implying that bonds in these markets have quite different yields. Notice that Municipal bonds have the lowest yields, which seems odd because Treasury bonds are safer investments. This relation is due to a tax-advantage that municipal bonds have. Interest from municipal bonds is not taxed at the federal, state, or local level whereas interest from Treasury bonds is only tax-exempt at the state and local levels.²⁷ Corporate bond yields, even for relative safe issuers, are higher than Treasury yields partly because of higher default risk and partly because of less liquidity. It's more difficult and costlier to trade corporate bonds than Treasury bonds.

The top right graph presents three different Treasury yield curves measured at different points in time. The graph highlights just how much interest rates can change. In early 1980, high inflation caused interest rates to rise above 15% on the short-end of the yield curve whereas interest rates following the great financial crisis in 2008 were frequently at levels

²⁷More precisely, municipal bond interest is not taxed at the state level if the investor lives in the state that issued the bond. Treasury bond interest is not taxed at the state or local level regardless of the state in which the investor lives.

below 2%. The graph also shows how the slope of the yield curve can vary over time from its typical upward slope in September of 1992 to flat in July 2019 to downward sloping or **inverted** in March 1980.

The bottom middle graph shows how yield curves can even vary within the same market, broadly defined. The graph shows three yield curves for municipal bonds with different credit ratings - A, AA, and AAA. The AAA yield curve is everywhere below the AA yield curve, consistent with AAA-rated bonds being less risky than AA-rated bonds. Interestingly, after seven years, A-rated bond yields are lower than AA-rated bond yields and after 13 years lower than AAA-rated bond yields despite being riskier than both. The most likely explanation for this pattern is higher liquidity for A-rated bonds relative to the others. There may be more trading activity and interest in A-rated bonds that has a larger impact on their yields than the difference in default risk, which is quite small.²⁸

7.10.1 What Does the Term Structure Mean?

Term structures across markets and time differ for a variety of reasons discussed above: default risk, inflation, liquidity. But, why does a specific term structure look the way it does? Why does it slope up? Down? There are several theories, but perhaps the most popular is the **expectations hypothesis**, which loosely says that the slope of the term structure is governed by expectations about future interest rates. An upward sloping term structure in which long-term interest rates are higher than short-term interest rates is indicative of investors expecting *future* interest rates to increase. A downward sloping term structure is indicative of lower future interest rates.

Because interest rates tend to decline during bad economic times, an inverted (downward sloping) yield curve is often viewed as a harbinger of a recession - an empirical fact first documented by Campbell Harvey in 1988.²⁹ However, the expectations hypothesis doesn't tell us exactly when interest rates will change nor does it guarantee that they will change in the direction we expect. Interest rates change every day as market participants change the supply and demand for bonds. Consequently, the term structure changes frequently and mostly in unpredictable ways.

²⁸While not exactly comparable, the average default rates for AAA-, AA-, and A-rated corporate, not municipal, bonds was 0.00%, 0.01%, and 0.05%, respectively.

²⁹Campbell Harvey, 1988, "The real term structure and consumption growth," *Journal of Financial Economics*, 22(2), 305-333.

7.10.2 The Fed and Monetary Policy

Let's end this chapter by discussing the **federal reserve system**, or **the Fed**, which plays a central role not just in influencing interest rates but in the economy at large. The Fed is the U.S. central bank, which consists of.³⁰

1. a board of governors or “the board” that oversees federal reserve banks and sets monetary policy,
2. 12 federal reserve banks that supervise and provide financial services to member banks in different regions of the country, and
3. the federal open market committee (FOMC) that buys and sells securities in financial markets to affect monetary policy.

, These three entities are also responsible for maintaining stability of the financial system, fostering payment and settlement system safety and efficiency, and promoting consumer protection and community development. While all of these responsibilities are important for the financial health of the country and its citizens, monetary policy receives the most attention and is arguably the most important tool that the Fed has at its disposal.

Monetary Policy

Monetary policy consists of the actions and communications of the Fed, which are centered around influencing the money supply - the total amount of cash, coins, and money in bank accounts in the economy. The Fed employs several tools to influence the money supply, all of which either directly or indirectly influence interest rates. By influencing interest rates, the Fed attempts to maximize employment and maintain stable prices (i.e., limiting inflation). This is the so-called **dual mandate of the Fed**.

How do interest rate changes affect employment and inflation? Interest rates are central to personal and business decisions. For example, beginning in March of 2022, the Fed began increasing interest rates to combat rising inflation. By increasing interest rates, people will prefer to save more and spend less to take advantage of higher returns on their savings. Through less spending by consumers, demand for goods and services will fall and with it the growth in prices, i.e., inflation. Higher interest rates also slows firms' spending through

³⁰For more details, see the federal reserve website at: <https://www.federalreserve.gov/default.htm>.

a higher cost of capital, r , that makes projects less valuable. Hence, inflation is slowed by increasing interest rates.

Because raising interest rates increases r in our Fundamental Value Relation, asset values will fall absent an offsetting increase in cash flows. So, a consequence of the Fed raising interest rates - **monetary tightening** - is a reduction in asset values. That's exactly what's happened since the Fed began its policy of raising interest rates. Between March and September of 2022, the S&P 500 index, a barometer of US stock markets, fell by over 17%. The iShares Core US Aggregate Bond ETF, a barometer of US bond markets, fell by nearly 12%.

It is tempting to blame the declines in asset values entirely on the increase in the discount rate, r , stemming from the Fed's increase in interest rates. However, the increase in interest rates also likely affects cash flow expectations in a negative way. Reducing consumer spending means lower profits for firms and lower cash flows. So, valuations get a double whammy - higher discount rates and lower expected cash flows.

Now consider when the Fed reduces interest rates - **monetary easing** - which occurred following the onset of the great financial crisis in 2008. People were incentivized to spend money as the return to savings (i.e., interest rate) fell. The increase in demand also increased prices, which was welcomed at the time because the economy was in a period of deflation - falling prices.

Now go back to our fundamental valuation formula. Reducing interest rates decreases r thereby increasing asset valuations absent an offsetting reduction in cash flows. In other words, reducing interest rates reduces the denominator in our Fundamental Value Relation which will lead to higher values unless there is a corresponding reduction in the numerators, i.e., the cash flows. In fact, the increased demand accompanying reductions in interest rates often works to increase cash flows and amplify the effect on valuations. From 2009 to 2014, stock and bond markets experienced impressive returns reflecting increasing valuations over time.³¹

³¹Academics have been studying the transmission mechanisms of monetary policy, or how monetary policy effects the economy, for decades. References include: Milton Friedman and Anna J. Schwartz, 1963, *A Monetary History of the United States, 1867-1960*, Princeton University Press; Christina Romer and David Romer, "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomics Annual*, 4, 121-170; and Ben S. Bernanke and Mark Gertler, 1995, "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," *Journal of Economic Perspectives*, 9(4), 27-48.

7.11 Key Ideas

We've really just scratched the surface when it comes to fixed income instruments. There are lots of different types of bonds and loans. From a financial perspective, they differ only in their implications for cash flows and discount rates.

- Bond prices are computed like those of any other asset, the sum of discounted future cash flows. The cash flows for bonds consists of interest and principal payments.
- The yield to maturity is the one discount rate such that when we sum the discounted bond cash flows, we recover the price of the bond. The yield is the internal rate of return of the bond and is quoted as an APR.
- The yield curve is a visual representation of the term structure of interest rates. For each term structure, there is a yield curve. The term structure of interest rates tells us the interest rates for loans of different lengths or terms.
- Realized returns are the returns we as investors actually earn. Expected returns are the returns we expect to earn on average. The two are rarely equal.
- A replicating portfolio is a collection of assets whose cash flows exactly match - in terms of timing and magnitude - the cash flows of another asset.
- The **principle of no arbitrage** ensures the relative prices of assets do not get too far out of whack for long. The act of taking advantage of relative mispricing by arbitrageurs corrects these mispricings.

This is *not* to say that market prices are never wrong - i.e., prices deviate from value. We can all think of examples that are difficult to rationalize with the discounted cash flows model we are using (e.g., “The dot com bubble,” the Game Stop stock run-up). However, inflated prices are often difficult to combat because the opposing competitive force is short selling. Our commentary on short selling shows that comes with a lot of risk, which imposes **limits of arbitrage**.³² Simply put, arbitrage or mispricing opportunities can be large and persist because costly realities limit the ability of investors to correct these prices.

- Bond risk is comprised of interest rate risk, default risk, and inflation risk (for nominal bonds). We measure interest rate risk with duration measures such as DV01 and

³²This term was coined by Andrei Shleifer and Robert Vishny in their 1997 *Journal of Finance* article by the same name.

modified duration. Default risk is comprised of two components: the probability of default and the loss given default.

- The Fed is the U.S. central bank, which is responsible for the supervision other banks in the U.S. and the execution of monetary policy. Guided by the dual mandate to keep inflation at 2% and the economy at full employment, monetary policy affects short-term interest rates and, by extension, broader financial and economic conditions through the discount rate, r , in our Fundamental Value Relation. This is why the Fed gets so much attention. When they change interest rates, they affect a lot of financial decision making - whether banks want to lend, whether companies want to invest or hire, whether people want to save or spend, and so on.

7.12 Technical Appendix

7.12.1 Forward Rates

When we agree *today* upon an interest rate for a loan that we will make in the *future*, that interest rate is called a **forward interest rate** or simply **forward rate**. Figure 7.31 illustrates spot and forward rates.

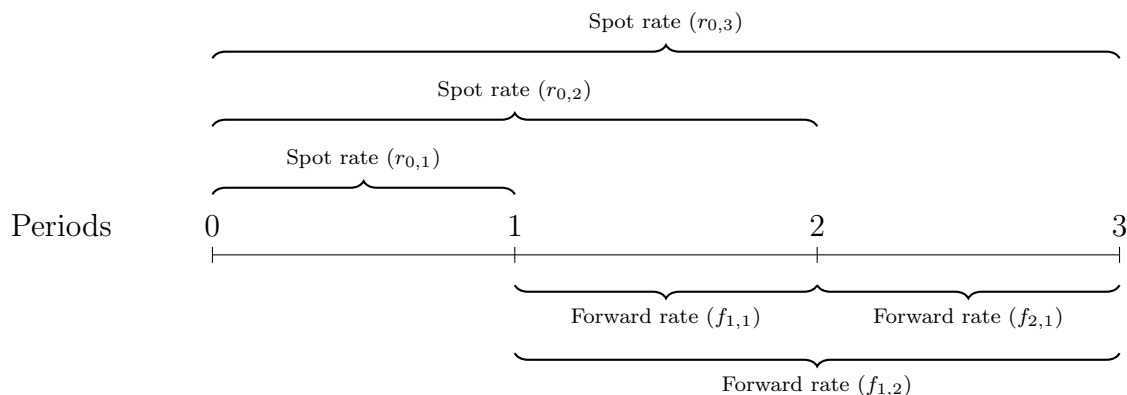


Figure 7.31: Spot and Forward Interest Rates

The spot rates $r_{0,1}$, $r_{0,2}$, and $r_{0,3}$ are the periodic interest rates for a 1-, 2-, and 3-period loan, respectively. Importantly, all of these loans begin *today*. The forward rates $f_{1,1}$, $f_{2,1}$, and $f_{1,2}$ are interest rates we agree upon today for loans that will start in the future. For example, $f_{1,2}$ is the interest rate for a loan that will start one period from today and last for one period, i.e., from period 1 to 2. Likewise, $f_{2,1}$ is the interest rate for a loan starting two periods from today and lasting for one period, i.e., from period 2 to 3. Finally, $f_{1,2}$ is the

interest rate for a loan starting one period from today and lasting for two periods, i.e., from period 1 to 3. The subscripts s, t correspond to the starting period and duration of the loan.

It is important to distinguish between forward rates and *future* spot rates. The former are known today, the latter are not. For example, the rates at which agree *today* to lend money in the future - $f_{1,1}, f_{1,2}, f_{2,1}$ - need not be the same as the spot rates that occur in the future.

There is an important relation between current spot rates and forward rates. Intuitively, it shouldn't matter how we lend or borrow money between two points in time as long as we know all of the terms today. For example, we could borrow money today for two periods at the current spot rate $r_{0,2}$. Or, we could borrow money for one period at the current spot rate, $r_{0,1}$, and agree *today* to borrow again for one more year at the one-year forward rate, one-year from now, $f_{1,1}$. Mathematically, this equivalence is as follows.

$$(1 + r_{0,2})^2 = (1 + r_{0,1}) \times (1 + f_{1,1})$$

This relation is visualized in figure 7.32.




Figure 7.32: Spot-Forward Rate Relation

Because the cost of these two strategies should be the same, the one-period forward rate is determined by the current one- and two-period spot rates.

$$f_{1,1} = \frac{(1 + r_{0,2})^2}{1 + r_{0,1}} - 1$$

(We just solved the previous equation for $f_{1,1}$.)

More generally, the cost of borrowing money at the T -period spot rate, $r_{0,T}$ must equal the cost of borrowing money for $s < T$ periods at the s -period spot rate, $r_{0,s}$ and then agreeing today to borrow for another $T - s$ periods at time s . That's a mouthful but what we're saying can be expressed mathematically as follows.

$$(1 + r_{0,T})^T = (1 + r_{0,s})^s \times (1 + f_{s,T-s})^{T-s} \quad \text{for } 0 < s \leq T \quad (7.9)$$

This expression is visualized in Figure 7.33.




Figure 7.33: Spot-Forward Rate Relation

Solving equation 7.9 for the forward rate produces a general forward rate relation.

$$f_{s,T-s} = \left[\frac{(1 + r_{0,T})^T}{(1 + r_{0,s})^s} \right]^{\frac{1}{T-s}} - 1 \quad (7.10)$$

Equation 7.10 tells us what the forward rate must be for a loan made in period s for $T - s$ periods in duration. If the forward rate is *not* equal to the right side of equation 7.10, then there is an arbitrage opportunity. See the problem set for an illustration of the arbitrage.

7.12.2 Reconciling Interest Rate Risk Measures

Mathematically, modified duration is defined as follows.

$$\text{Modified Duration} = -\frac{1}{P_t} \frac{\partial P_t}{\partial y} \quad (7.11)$$

This expression shows that the negative of modified duration tells us the relative change in bond prices, $\partial P_t / P_t$, for a small change in yields ∂y .

Multiplying both sides of equation 7.11 by -1 and ∂y yields the following expression.

$$\frac{\partial P_t}{P_t} = -\text{Modified Duration} \times \partial y$$

Now plug in 0.01 for ∂y and multiply *both* sides by 100 to convert the units to percent.

$$\underbrace{\frac{\partial P_t}{P_t} \times 100}_{\% \text{ Price Change}} = -\text{Modified Duration}$$

So, the negative of modified duration is the percentage change in price for a one percent change in yield. For example, if modified duration is -5.68 , then a 1% change in yield leads to a 5.68% change in the bond price.

Now let's relate modified duration to DV01. Multiply both sides of equation 7.11 by the price, P_t and 0.0001.

$$-\text{Modified Duration} \times P_t \times 0.0001 = \underbrace{\frac{\partial P_t}{\partial y} \times 0.0001}_{DV01}$$

The right side of this equation is DV01, the dollar value change in the bond price for a one basis point change in yield. This equals the negative of the modified duration times the bond price times one basis point.

7.12.3 Taxes on Zero Coupon Bonds

Earnings on all virtually investments are taxed in the U.S. Interest income is taxed as ordinary income. The specific tax rate depends on the tax bracket in which we fall - higher income, higher taxes. Capital gains are taxed similarly, but as of 2021 they are taxed at a lower rate than interest income. Consider a 10-year, zero-coupon bond with a face value of \$1 million, 4% semi-annual compounded APR, and price of \$672,971.33. Because the bond price is below the face value at issuance, the bond is referred to as an **original-issue discount** or **OID**.

The lack of interest payments could create a tax preference for zero coupon bonds if they are only taxed at the capital gains rate, which is lower than the tax on ordinary income for many investors. (As of 2022, the top federal tax rate on ordinary income was 37%, on capital gains it was 20%.) However, the tax authority recognizes that the price discount is a way for the bond to implicitly earn interest over its term. Thus, OID bonds are taxed as if the bond earned interest each year, even though the investor doesn't receive any interest.

Take a 10-year zero coupon bond. One year from today, if interest rates are unchanged, the bond will be priced at

$$P_1 = \frac{1,000,000}{(1 + 0.04/2)^{18}} = \$700,159.38.$$

The tax authority deems the price change $700,159.38 - 672,971.33 = \$27,188.05$ to be earned interest subject to ordinary income tax. In other words, the implied interest on an OID bond is computed from the annual price difference *assuming* the yield on the bond is unchanged.

Now imagine that interest rates change from 4% to 3% over the year. The price one year from today would be

$$P_1 = \frac{1,000,000}{(1 + 0.03/2)^{18}} = \$764,911.59.$$

If we sell the bond, the capital gains would be $764,911.59 - 700,159.38 = \$64,752.21$, which would be subject to the capital gains tax rate. If the bond is not sold, the capital gains are not realized and there is no tax implication. Regardless, taxes on the \$27,188.05 of interest income must be paid at the end of the year.

The lesson here is that we have to pay taxes on zero coupon bonds over time even if we don't sell them.

7.12.4 Taxes on Coupon Bonds

Taxes on coupon bond investments work similar to those on zero coupon investments. The interest is taxed as ordinary income and any gains or losses from the sale of the bond are taxed as capital gains or losses after adjusting for any original issue discount (OID), as discussed in the previous section. To illustrate, consider the following example.

Imagine we purchase a two-year treasury note with a 5% coupon that was issued on December 31, 2021 for \$108.38, and we sell it for \$102.27 18 months later. Also assume that our income and filing status (e.g., single, married filing jointly, head of household) puts us in the 32% tax bracket. The cash flows are detailed in figure 7.34.

Date	12/31/2021	6/30/2022	12/31/2022	6/30/2023
Cash flows	-108.38	2.50	2.50	104.77
Taxes		0.80	0.80	-1.16
After-tax cash flows	-108.38	1.70	1.70	105.93

Figure 7.34: Two-Year Treasury Note Cash Flows

Each coupon payment is taxed as ordinary income at 32%, implying $0.32 \times 2.50 = \$0.80$ in taxes. When we sell the bond on 6/30/2023, two things happen. The coupon gets taxed at 32% and we owe \$0.80. However, we experience a capital loss of $108.38 - 102.27 = \$6.11$, which we can deduct from our taxable income.³³ The capital loss reduces our taxes by $0.32 \times 6.11 = \$1.96$. The net result is that when we sell the bond, our taxes are reduced by $1.96 - 0.80 = \$1.16$. The pre-tax periodic return on our investment is 0.4367%, 0.8753% on an annual basis. The after-tax periodic return is 0.2951%, 0.5912% on an annual basis. (Periodic returns are simply the internal rates of return on the before- and after-tax cash flows of the bond.)

³³As of 2022, there was a \$3,000 annual limit on deducting capital losses from your income. Any losses beyond \$3,000 could be carried forward to reduce future taxable income.

7.12.5 Proofs

Inverse Relation Between Bond Prices and Yields

The derivative of our Fundamental Value Relation with respect to yield, after replacing *Value* with *Price* and the discount rate r with the bond yield y , produces the following expression, where k is the compounding frequency.

$$\begin{aligned}\frac{\partial P_t}{\partial y} &= \frac{1}{k} \left[-\frac{CF_{t+1}}{(1+y/k)^2} - 2\frac{CF_{t+2}}{(1+y/k)^3} - \dots - T\frac{CF_T}{(1+y/k)^{T+1}} \right] \\ &= -\underbrace{\frac{1}{(1+y/k)}}_{>0} \underbrace{\left[\sum_{s=1}^T \frac{s}{k} \times \frac{CF_{t+s}}{(1+y/k)^s} \right]}_{>0}\end{aligned}$$

The derivative equals the negative of the product of two positive terms, so the derivative is negative, implying a negative relation between yields and prices of bonds without embedded options.

Bond Prices, Yields, and Coupon Rates

The price of a bond can be expressed in terms of its yield as follows.

$$Price_t = \frac{c/k \times Face}{(1+y/k)} + \frac{c/k \times Face}{(1+y/k)^2} + \dots + \frac{c/k \times Face}{(1+y/k)^T} + \frac{Face}{(1+y/k)^T}$$

The coupon rate, c , and yield, y are both expressed as APRs with compounding frequency k . The discounted coupon payments are a geometric series whose partial sum, S_T , scaled by $(1+y/k)^{-1}$ is equal to

$$\begin{aligned}\frac{1}{1+y/k} S_T &= \frac{c/k \times F}{1+y/k} \left[\frac{1}{(1+y/k)} + \frac{1}{(1+y/k)^2} + \dots + \frac{1}{(1+y/k)^T} \right] \\ &= \frac{c/k \times F}{1+y/k} \left[\frac{1 - (1+y/k)^{-T}}{1 - (1+y/k)^{-1}} \right] \\ &= c/k \times F \left[\frac{(1+y/k)^T - 1}{y/k(1+y/k)^T} \right].\end{aligned}$$

Therefore, the price of a coupon bond can be expressed as follows.

$$Price_t = \frac{c/k}{y/k} \times F \left[1 - \frac{1}{(1+y/k)^T} \right] + \frac{F}{(1+y/k)^T} = \frac{c}{y} \times F$$

When the coupon rate is equal to the yield, $c = y$, the bond price is equal to the face value, $Price_t = F$ - the bond is priced at par. When $c < y$, the bond price will be less than the face value, $Price_t < F$, and the bond is priced at a discount. When $c > y$, the bond price will be greater than the face value, $Price_t > F$, and the bond is priced at a premium.

Zero Coupon Bond Interest Rate Sensitivity

The price of a t -year zero coupon bond rises and (falls) by about t -percent for each one percent decrease (increase) in interest rates.

$$P_t = \frac{CF_T}{(1+r)^{(T-t)}}$$

Let's compute the semi-elasticity of the price with respect to the interest rate.

$$\begin{aligned} \frac{1}{P_t} \frac{\partial P_t}{\partial r} &= -\frac{1}{P_t} \times \frac{T-t}{1+r} \times \frac{CF_T}{(1+r)^{T-t}} \\ &= -\frac{T}{1+r} \\ &\approx -T \text{ for small } r \end{aligned}$$

Macaulay Duration and the Derivative of the Bond Price

Start with the basic pricing formula using the annual yield on the bond.

$$P_t = \frac{CF_{t+1}}{(1+y/k)} + \frac{CF_{t+2}}{(1+y/k)^2} + \dots + \frac{CF_{t+T}}{(1+y/k)^T}$$

Take the derivative with respect to the yield.

$$\begin{aligned} \frac{\partial P_t}{\partial y} &= -\frac{1}{k} \frac{CF_{t+1}}{(1+y/k)^2} - \frac{2}{k} \frac{CF_{t+2}}{(1+y/k)^3} - \dots - \frac{T}{k} \frac{CF_{t+T}}{(1+y/k)^{T+1}} \\ &= -\frac{1}{1+y/k} \left[\frac{1}{k} \frac{CF_{t+1}}{(1+y/k)} + \frac{2}{k} \frac{CF_{t+2}}{(1+y/k)^2} + \dots + \frac{T}{k} \frac{CF_{t+T}}{(1+y/k)^T} \right] \end{aligned}$$

Multiplying by the inverse of the price produces the expression for modified duration.

$$\begin{aligned} \frac{1}{P_t} \frac{\partial P_t}{\partial y} &= -\frac{1}{1+y/k} \left[\frac{1}{k} \left(\frac{PV(CF_{t+1})}{P_t} \right) + \frac{2}{k} \left(\frac{PV(CF_{t+2})}{P_t} \right) + \dots + \frac{T}{k} \left(\frac{PV(CF_{t+T})}{P_t} \right) \right] \\ &= -\frac{1}{1+y/k} \underbrace{\sum_{s=1}^T \underbrace{\left(\frac{s}{k} \right)}_{\text{Year of } s^{\text{th}} \text{ cash flow}} \times \underbrace{\left(\frac{PV(CF_{t+s})}{P_t} \right)}_{\text{Weight of } s^{\text{th}} \text{ cash flow}}}_{\text{Macaulay Duration}} \end{aligned}$$

where $PV(CF_s) = CF_s/(1+y/k)^s$ is the present value of the cash flow arriving in period s .

7.13 Problems

7.1 (*Conceptual*) Why do bond prices decline when interest rates rise?

7.2 (*Zero coupon bond valuation*) Find the price today of a Treasury STRIP that matures in five years and has a face value of \$100 if the APR is 7.5% with semi-annual compounding.

7.3 (*Zero coupon bond yield and return*) Bob has decided to purchase a Treasury strip maturing three years from today for \$87.54 per \$100 of par value.

Using this information, answer the following questions.

- What is the annual yield-to-maturity of the bond?
- If one year from today the two-year yield to maturity is 6%, what will be the price of the bond?
- If Bob sells the bond one year from today, what will be his realized return?
- If instead bob decides to hold his bond until maturity, what will his annualized return be?

7.4 (*T-bill pricing*) Today is June 12, 2023 and we are considering purchasing a T-bill maturing on September 21, 2023. The bill's bond-equivalent yield (i.e., semi-annual compounded APR) is 5.3064%. What is a fair price for the bond?

7.5 (*T-bill trading*) Today is June 12, 2023 and we are given the following information about a Treasury bill.

Security Type	Coupon Rate (%)	Maturity Date	Bid	Ask
Bill	0.00	2/22/24	96.543333	96.589375

Using this information answer the following questions.

- How many days remain until maturity?
- How many years remain until maturity?
- How many semi-annual periods remain until maturity?
- At what price can we buy the bond? (Be sure to specify how much face value of the bond you can buy at that price.)
- At what price can we sell the bond? (Be sure to specify how much face value of the bond you can sell at that price.)
- What is the yield to maturity at the ask price?

- g. If we buy the bond today, what will our realized return be assuming we hold the bond to maturity and the U.S. government doesn't default? What is our annualized return?

7.6 (*T-bill trading*) Today is June 12, 2023 and we are given the following information about a Treasury bill.

Security Type	Coupon Rate (%)	Maturity Date	Bid	Ask
Bill	0.00	10/5/2023	98.348472	98.354861

Using this information answer the following questions.

- What is the bid-ask spread of the bond?
- How much would it cost to purchase \$5,000,000 of par value?
- How much interest, expressed in dollars, would we earn if we purchased \$5,000,000 of par value.
- What return does this bond offer, assuming the government does not default and we hold the bond to maturity? What is the annualized return?

7.7 (*Treasury STRIP quotes*) Today is June 13, 2023 and we are given the following information about a Treasury STRIP.

Security Term	Maturity Date	Bid/Ask Price	Bid/Ask YTM (%)	Bid/Ask Quantity
30-year	8/15/2049	35.02/36.77	4.05/3.86	3000/3000

The quantity quotes are the number of bonds, where each bond has \$1,000 par value, available for purchase or sale at the quoted prices.

Using this information answer the following questions.

- At what price can we purchase this bond? What is the corresponding yield to maturity?
- At what price can we sell the bond? What is the corresponding yield to maturity?
- On what date(s) does the bond pay its owner?
- If we purchase \$5,800 of face value of this STRIP, how much money will we receive when the bond matures?

- e. If we hold the bond to maturity and the government doesn't default, what will our realized return be from today to maturity? What will our annualized return be?

7.8 (*Treasury STRIP comparison*) Today is June 13, 2023 and we are given the following information about two Treasury STRIPs.

Bond	Security Term	Maturity Date	Bid/Ask Price	Bid/Ask YTM (%)	Bid/Ask Quantity
A	30-year	5/15/2028	81.89/82.12	4.10/4.04	2,000/5,000
B	30-year	2/15/2038	54.93/55.26	4.13/4.08	2,000/2,000

The quantity quotes are the number of bonds, where each bond has \$1,000 par value, available for purchase or sale at the quoted prices.

Using this information answer the following questions.

- Is bond B a better investment than bond A because it's yield-to-maturity is higher? Explain.
- How much would it cost to buy 1,500 A bonds?
- How much would we receive for selling 570 B bonds?
- How many A bonds could we purchase with \$475,000, assuming we cannot buy fractional units of a bond (i.e., no face values less than \$1,000)? How much would receive at maturity from our bonds if we did purchase them?
- How many B bonds could we purchase with \$475,000, assuming we cannot buy fractional units of a bond (i.e., no face values less than \$1,000)? How much would receive at maturity from our bonds if we did purchase them?
- (*Challenging*) Using the results from the previous two questions, imagine taking the proceeds from investing in bond A when it matures on 5/15/2028 and reinvesting in another Treasury STRIP that matures on 2/15/2038. What yield would this bond have to promise in order to ensure the proceeds from this second investment would match those from just investing in bond B today? What is the price of this bond per \$100 of par value?

7.9 (*Coupon bond valuation*) You are offered an 8% coupon bond with a face value of \$1,000,000. The bond has three years to maturity, the coupons are paid semi-annually, and the yield-to-maturity is 10%

Using this information, answer the following questions.

- a. How much are the coupon payments?
- b. What is the present value of the coupon payments?
- c. How much is the principal amount of the bond?
- d. What is the present value of the principal?
- e. What is the present value of the bond? Assuming this value is equal to the price, is the bond priced at a discount, premium, or at par?

7.10 (*J&J coupon bond yield*) Johnson and Johnson Corp. is preparing to issue a semi-annual coupon paying bond. It would like to see the bond yield equal 5%, indicative of the current interest rate environment. What must the annual coupon rate be to ensure the bond is priced at par?

7.11 (*IBM coupon bond valuation*) IBM has just issued a 20-year semi-annual coupon bond with a 6.25% coupon rate. The yield on the bond is 4.86%.

Using this information, answer the following questions assuming a par value of \$100.

- a. What is the present value of the coupon payments?
- b. The coupons correspond to what type of cash flow stream?
- c. What is the present value of the bond face value?
- d. What is the present value of the bond? Assuming this value is equal to the price, is the bond priced at a discount, premium, or at par?
- e. At what coupon rate will the bond be priced at par?

7.12 (*T-note yield*) (*Challenging*) On June 1 2023, the U.S. Treasury 4.25% of 9/30/2024 had an ask price of 98.984375 per \$100 of par value. What is the current yield of this bond assuming the next two coupon dates are 9/30/w023 and 3/31/2024? What is the annualized expected return assuming the government does not default and we hold the bond to maturity?

7.13 (*T-bond quotes*) Today is June 12, 2023 and we are given the following information about a Treasury bond.

Security Type	Coupon Rate (%)	Dated Date	Maturity Date	Bid	Ask
Bond	2.50	2/15/2016	2/15/2046	76.34375	76.43750

Using this information, answer the following questions.

- What is the original term of the T-bond?
- When is the next coupon payment?
- How many coupon payments remain?
- What is the bid yield-to-maturity if today were 2/15/2023?
- What is the ask yield-to-maturity if today were 2/15/2023?
- (*Challenging*) What is the bid yield-to-maturity as of June 12, 2023?
- (*Challenging*) What is the ask yield-to-maturity as of June 12, 2023?

7.14 (*Valuation, yield calculation, and arbitrage strategy*) Consider the following market data on four risk-free bonds.

Bond	Par Value	Coupon Rate (%)	Maturity (Years)	Yield (%)	Price
A	100	0	1	?	95.2381
B	100	0	2	?	90.7030
C	100	0	3	?	83.9620
D	1000	10%	3	5.00	?

The coupon rate and yield are APRs. Coupons and interest are paid and compounded annually.

Using this information, answer the following questions.

- What are the values for each “?” in the table?
- Is there an arbitrage opportunity? If so, how would you take advantage of it? Be clear to identify the positions and corresponding cash flows.

7.15 (*Realized return, yield, interest rate changes and valuation*) Your recently received \$1.5 million in cash from the sale of your home. You would like to park this cash somewhere safe and liquid. You decide to put your money into a 6-month Treasury bill whose current price is \$98.50 per \$100 of par value.

Using this information, answer the following questions.

- Assuming the federal government does not default on the bond, what will be the realized return on the bond when it matures?
- What is the yield to maturity of the bond?

- c. What is the annualized return on of the bond? What does this calculation assume the investor does to earn this return?
- d. If you invest all \$1.5 million in the 6-month Treasury, how much money will you have at the end of six months?
- e. If immediately after purchasing your bond the bond yield increases by 75 basis points because of a change in monetary policy, what is the new price of your bond? If you hold the bond to maturity, what impact will this interest rate change have on your realized return?

7.16 (*Bond portfolios, bond ladders*) When filing his taxes every April, Ryan sets aside money to cover his quarterly estimated tax payments for the coming year. This year those payments will be \$40,000 each.

These payments are due in April (immediately), July, October, and January. To ensure his money is both safe and earns a market rate of return, he wants to invest his money in Treasury securities by constructing a bond ladder. A bond ladder is a portfolio of bonds with staggered maturities. In this instance, Ryan wants to purchase 3-, 6-, and 9-month T-bills to coincide with his upcoming tax payments.

The short end of the current Treasury yield curve is detailed in the following table. (Remember that Treasury yields and their coupon rates are semi-annual compounded and quoted as APRs.)

Term	Yield (%)
3 Mo	4.90
6 Mo	4.79
9 Mo	4.60

Using this information, answer the following questions.

- a. What is the present value of Ryan's future tax obligations?
- b. What is the price of each T-bill?
- c. T-bills can be purchased in \$100 increments or "units." How many units of the 3-, 6-, and 9- month T-bill must Ryan purchase today to ensure he can make each quarterly payment? What is the total cost of each purchase? Put differently, how much money must he spend on 3-, 6-, and 9- month T-bills?

7.17 (*Bond portfolios, bond ladders*) Continuing from the previous problem, Ryan is considering an alternative strategy in which he buys a two-year Treasury note with a 5.25% coupon. The zero coupon Treasury yield curve is presented in the table below.

Term	3 Mo	6 Mo	9 Mo	1 Yr	1.25 Yr	1.5 Yr	1.75 Yr	2 Yr
Yield (%)	4.90	4.79	4.60	4.88	4.92	5.02	5.31	5.75

Using this information, answer the following questions.

- What is the price of the 2-year T-note per \$100 of face value?
- What is the yield-to-maturity of the 2-year T-note?
- Like T-bills, T-notes can be purchased in \$100 increments or “units.” How many units of the 2-year T-note must Ryan purchase today to ensure he can make each quarterly payment assuming interest rates do not change? What is the total cost of his purchase?
- (Challenging)* What is the price of the 2-year T-note per \$100 of face value 3 months from today if interest rates don’t change? Explain what happened to the price?
- (Challenging)* How many bond units must Ryan sell every three months to meet his tax obligations assuming interest rates don’t change? Will there be any excess capital after the last payment, or by how much will he fall short?
- (Challenging)* Explain the relation between any shortfall or excess capital Ryan will have 9 months from now after his last tax payment and the slope of the yield curve.

7.18 (*Interest rate changes and bond valuation, bonds as a savings strategy*) You have to pay \$20,000 for your child’s college at the end of the next two years. You can invest your money today in two different types of investments:

- one-year zero-coupon bonds, and
- four-year zero coupon bonds.

Assume the yield curve is flat at 8% per annum. Interest is compounded semi-annually for both bonds.

Using this information, answer the following questions.

- How much money must you put aside today assuming that interest rates stay unchanged?
- What are the prices of the one-year and four-year zero coupon bonds per \$100 of par value?

- c. (*Challenging*) Suppose you invest the money you need to set aside today - your answer to the first question - in 4-year bonds. What is the par value of the bonds you need to sell one year from today to meet the first payment? If interest rates remain unchanged, show that you will exactly meet your obligations over the next two years.
- d. (*Challenging*) Consider the previous question in light of two different scenarios. In the first, interest rates permanently increase to 9% and in the second, interest rates permanently decrease to 7%. Assume both changes take place immediately after buying the bonds. For each scenario, answer the following questions. What is the par value of the bonds you need to sell one year from today to meet the first payment? After the second payment, how much extra money will have left over, or will you have enough to make the second payment?

7.19 (*Hedging interest rate risk*) You have a liability (amount owed) of \$100 million due five years from today. You want to invest today in securities to protect against any future interest rate risk. You can invest in 3-year and 7-year, risk-free zero-coupon bonds.

Using this information, answer the following questions.

- a. How much should you invest in each bond if we assume that the term structure is flat at 5%?
- b. Show exactly how the hedge constructed in the previous problem works by computing the value of the positions you came up with in the previous problem after five years. What happens to these values if interest rates change to 3% or 7% immediately after establishing your asset position?

7.20 (*Spot and forward rates, bond valuation with the yield curve*) The one-year annually compounded spot rate is 2%. The one-year annually compounded forward rates one-year and two-years hence are 4% and 5%, respectively.

Using this information, answer the following questions.

- a. What are the two- and three-year spot rates?
- b. Using the spot rate yield curve you constructed in the previous question, what is the price of a three-year annual coupon paying bond with \$1,000 par value and 8% annual coupon rate? What is the yield-to-maturity of this bond? Explain the relation between the bond yield, the coupon rate, and whether the bond price is above or below the par value.

7.21 (*Coupon bond yield*) A 3-year Treasury note is currently priced at \$94.67 per \$100 of par value. The note has a 4% coupon rate, with semi-annual coupons. What is the yield-to-maturity expressed as an APR?

7.22 (*Bond yields and valuation, before- and after-tax realized returns*) You are exploring the following two bonds.

- (a) Bond A matures in two years and carries a 10% annual coupon rate.
- (b) Bond B matures in two years and pays no coupons.

Both bonds have a 10% yield to maturity. The income tax rate is 36%, and the capital gains tax rate is 15%.

Using this information, answer the following questions.

- a. What is the price of each bond per \$100 of par value?
- b. What are the before- and after-tax holding period returns of each bond one year from today assuming that they bond yields do not change? (You should compute four numbers: one for each bond on a before- and after-tax basis.)
- c. What are the before- and after-tax holding period returns of each bond one year from today assuming that they bond yields decrease by 1%, from 10% to 9%? (You should compute four numbers: one for each bond on a before- and after-tax basis.)

7.23 (*Bond yields and valuation, realized returns, interest rate risk*) Alex is leaving for college next year and would like to invest his college savings in a fixed income product. He has three choices, all of which are maturing five years from today.

- (a) a Treasury strip currently trading at \$78.1198 per \$100 of par value.
- (b) a 5-year Treasury note with a 5% semi-annual coupon and a 5% yield to maturity.
- (c) a 10-year Treasury note and a 7% semi-annual coupon and a 5% yield to maturity.

Using this information and assuming semi-annual compounding for all of the bonds, answer the following questions.

- a. What is the yield to maturity of the Treasury strip?
- b. What are the prices of the Treasury notes - the 5-year and the 10-year - per \$100 of par value?

- c. Assuming interest rates don't change, what will be the price of each bond one year from today when Alex cashes out to pay for college? What will be his before tax holding period return for each bond?
- d. If instead interest rates increase by 1.0% over the next year - think a parallel shift of the yield curve - what will be the price of each bond, and what will Alex's before-tax returns to each bond be?
- e. What do your results in the previous problem reveal when it comes to the potential risk Alex is taking by investing his money in these safe securities? Support your argument by computing the modified duration of each bond as of today.

7.24 (*Coupon bond valuation*) You've been asked to value following delayed coupon bond per \$100 of par value. The bond matures in 15 years from today and has a 7% annual coupon rate. However, instead of the first coupon being paid one year from today, the first coupon is paid six years from today and each subsequent coupon occurs every year thereafter until and including at maturity. Assume a 10% APR with annual compounding.

7.25 (*Perpetual bond valuation*) You've been asked to value a perpetuity that pays \$1,000 at the end of odd years and \$1,500 at the end of even years. Compute the price of this perpetuity if the current annual yield is 10%.

7.26 (*Zero coupon bond valuation, arbitrage strategy*) You are given the prices for two zero coupon bonds maturing in one year.

Bond	Price	Par
A	90	100
B	285	300

Compute the implied yield-to-maturity for each bond. Is there an arbitrage opportunity? If so, detail the positions you would take, and corresponding cash flows, to take advantage of the opportunity. What real world considerations (often referred to as **market imperfections** or **financial frictions**) might prevent us from taking advantage of any arbitrage opportunity?

7.27 (*Spot interest rates, bond yields*) You are given the following information on three bonds, all of which have par values of \$100, annual coupon payments where if relevant, and annual compounding.

Bond	Coupon rate (%)	Maturity (Years)	Price (\$)
A	0	1	90
B	10	2	102.50
C	0	2	?

What are the implied one- and two-year spot rates? Is the yield-to-maturity on bond B greater than, less than, or equal to the coupon rate? What is the price of bond C?

7.28 (*Spot interest rates, bond yields*) What is the price today of a 10-year Treasury note with a face value of \$1 million, 10% annual coupon rate, and semi-annual coupons if the annual yield on the bond is 8%.

What is the price of the same bond with two years from today (i.e., just after the fourth coupon payment) if interest rates have not changed?

7.29 (*Measuring default risk*) Altman's Z-score is popular measure characterizing the default risk of a company. It is defined as follows.

$$Z = 3.10 \frac{EBIT}{TotalAssets} + 1.00 \frac{Sales}{TotalAssets} + 0.42 \frac{BookEquity}{TotalLiabilities} + 0.85 \frac{RetainedEarnings}{TotalAssets} + 0.72 \frac{WorkingCapital}{TotalAssets}$$

Does a higher Z-score correspond to a firm with a higher or lower likelihood of financial distress? Explain your answer.

7.30 (*After-tax bond yield comparisons*) Consider the Treasury and Municipal bond yield curves in figure ???. Assume the following.

- Taxes are paid once a year.
- The federal tax rate is 32%.
- One-year bonds are zero-coupon and issued at a discount.
- Bonds with maturities greater than one year make annual coupon payments and are issued at par.
- Interest is compounded annually.

Using this information, answer the following questions.

- a. What are the prices of the one-year Treasury (i.e., T-bill) and municipal bonds?

- b. What are the after tax-yields on the one-year Treasury and municipal bonds? Compute these after-tax yields in two ways. First compute the internal rate of return on the bond. Second, compute

$$\text{After-tax bond yield} = \text{Pre-tax bond yield} \times (1 - \text{Tax rate})$$

How do these two estimates compare? Explain any difference.

- c. How do the after-tax yields on the Treasury and municipal bond computed in the previous question compare? Explain any difference between these yields?
- d. Answer the previous questions again using the thirty-year Treasury and municipal bonds? Do these longer-term bonds come with any additional risks?

7.31 (*After-tax bond yield comparisons*) Consider the Treasury and Municipal bond yield curves in figure ??.

- Taxes are paid once a year.
- The federal tax rate is 32%.
- One-year bonds are zero-coupon and issued at a discount.
- Bonds with maturities greater than one year make semi-annual coupon payments and are issued at par.
- Interest is compounded semi-annually.

Using this information, answer the following questions.

- a. What are the prices of the one-year Treasury (i.e., T-bill) and municipal bonds?
- b. What are the after tax-yields on the one-year Treasury and municipal bonds? Compute these after-tax yields in two ways. First compute the internal rate of return on the bond. Second, compute

$$\text{After-tax bond yield} = \text{Pre-tax bond yield} \times (1 - \text{Tax rate})$$

How do these two estimates compare? Explain any difference.

- c. How do the after-tax yields on the Treasury and municipal bond computed in the previous question compare? Explain any difference between these yields?
- d. Answer the previous questions again using the two-year Treasury and municipal bonds? Do these longer-term bonds come with any additional risks?

7.32 (*Bank risk*) West Bank is a commercial bank located in the California and Arizona. It's market value balance sheet is presented below. (A market value balance sheet operates just like an accounting balance sheet except the values are market values as opposed to historical costs.) Dollar values are in billions. Macaulay duration is in parentheses.

Assets	Liabilities & Shareholders Equity
Treasurys (Duration = 7.8) 46	Debt (duration 2.5 years) 38
	Equity 8

The current yield curve is flat 2.0%.

Using this information, answer the following questions.

- a. What is the implied equity duration?
- b. If the Fed announces it will increase interest rates by 0.75%, how will the market value balance sheet change? Estimate the new values for each line item assuming the entire yield curve shifts up by 1%.

7.33 (*Forward rate arbitrage*) (*Challenging*) Annually compounded yields for 1-, 2-, and 3-year zero coupon bonds are in the following table.

Term (Years)	1	2	3
Yield (%)	4.5	5.0	6.2

Using this information, answer the following questions.

- a. What are the implied forward rates: $f_{1,1}$, $f_{2,1}$, $f_{1,2}$?
- b. The current market forward rate, $f_{1,1}$, is 6%. Is there an arbitrage opportunity? If so, how would you take advantage of it using one- and two-year bonds?
- c. Reconsider your answer to the previous question if the supply of two-year bonds was such that it is prohibitively costly to borrow them.

Chapter 8

Investing: Stocks

Fundamental Value Relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

This chapter answers the following questions.

- What are stocks?
- How do we buy and sell stocks?
- How do we make or lose money investing in stocks?
- How can we value stocks?
- What are the risks of investing in stocks?

In the process of answering these questions, we'll discuss

- different types of stock - e.g., common and preferred - and how they differ,
- different types of trades and the associated risks,
- what drives stock returns - both realized and expected,
- several different models of stock valuation - all of which are applications of the Fundamental Value Relation but with different assumptions about how to measure cash flows, and

- apply our Fundamental Value Relation to answer several specific questions including:
 - What is the current value of Microsoft and how does it compare to its market price?
 - What does the price-to-earnings ratio (P/E ratio) and price-to-earnings-to-growth ratio (PEG ratio) tell us about a company's price and relative valuation?
 - How can we identify whether a stock is over- or under-valued by the market?
 - How can we estimate an expected stock return and the risk of investing in a stock?

8.1 What is Stock?

Stock, or **equity**, is a claim to the cash flows generated by a firm, just like a loan. Unlike a loan, most stock is not promised any cash flows, and most stock comes with the right to vote on important corporate actions, like mergers and acquisitions or changes to the company's bylaws - the rules governing the behavior of the company. Individuals owning stock are called **shareholders** because stock is sold in units called **shares**. Shareholders own the company. That is, shareholders own all the assets and are responsible for the company's decisions, in the same way a homeowner owns their home and is responsible for all the actions concerning the home even if they borrowed money to purchase it.

8.1.1 Decision Making at the Company

Assuming each share is entitled to one vote, whomever owns the majority of shares has the majority vote and is entitled to make the decisions. For example, if a company has 100 shares outstanding, and Joe owns 20 shares and Mary 80, then Mary has the majority vote and is able to make all corporate decisions. If instead there are three shareholders - Joe, Mary, and Andreea - each owing one third of the company, then two of the shareholders would have to agree to any changes because a majority is needed. In some cases, a simple majority - 51% of the vote - is insufficient for some decisions. A supermajority - typically 67% - 90% - may be required for major corporate actions such as mergers, acquisitions, divestitures and changes to the board compensation.

Depending on the size of the company, corporate decisions may be made by the shareholders or delegated to others. Small, privately held firms are often owned and operated by the same person or people. Think of a local restaurant or convenience store. An individual or small group of people both own (i.e., are shareholders) and operate the business. For a

large publicly traded company such as Apple Inc., which has tens of thousands of shareholders, this management model is infeasible. Coordinating all these different owners to make decisions would not only be extremely costly but also a bad idea. Most shareholders are not experts in technology. Thus decision making at Apple is carried out by its employees, some of whom happen to be shareholders.

When shareholders and managers are different people, there is the potential for **conflicts of interest**. What's in the best interest of management - flying on private jets, hiring friends, investing in "pet projects," etc. - may not be in the best interest of shareholders. That is to say there are **incentive conflicts** between owners and managers. This separation of ownership (shareholders) and control (managers) is why many companies have a **board of directors** whose job is to ensure managers are acting in the interests of shareholders. Likewise, **shareholder meetings** are often required in which any registered shareholder of the company is allowed to attend and, in many cases, speak to voice concerns. While imperfect, these mechanisms play an important role in **corporate governance**.

8.1.2 Shareholder Cash Flows

Whereas lenders receive cash flows in the form of interest and principal, shareholders receive cash flows in the form of **dividends** - paid in cash and occasionally additional shares - and **stock buybacks** or **repurchases**. Dividends tend to be consistent payouts, much like a coupon payment on a bond, and typically occur on a quarterly basis. Dividends also tend to increase over time, albeit modestly, consistent with shareholder expectations. However, dividends can also be reduced or "cut", and even omitted when companies have financial difficulties because of a downturn in business, for example. Figure 8.1 illustrates two very different dividend trajectories for Coca-Cola (KO) and General Electric (GE).

Stock repurchase programs are initiatives in which the company buys back shares from its shareholders. Stock buybacks afford companies more flexibility in the timing and magnitude of their distributions to shareholders because, unlike dividends, buyback programs do not come with any expectations of consistent distributions to shareholders. Figure 8.2 highlights this flexibility by showing the repurchase activity of Proctor and Gamble (PG) and Exxon Mobil (XOM), two firms with some of the largest repurchasing activity between 2010 and 2022.

A key difference in the cash flows to lenders and shareholders is that the company is contractually required to make interest and principal payments or risk being thrown into bankruptcy by lenders. Dividend and stock buybacks are at the discretion of the company




Figure 8.1: Dividend per Share




Figure 8.2: Share Repurchases

and many companies, particularly younger, smaller companies, do neither for some time. Though, as we'll see, there must be some expectation that the company will return money to its shareholders or else they should never invest in the first place.

Another difference between lenders and stockholders is that the latter only get paid *after* all of the lenders have been paid. Stockholders are lower in the **priority structure** of companies' liabilities and, as such, are referred to as **residual claimants**. They get what's left over, if anything, after creditors are paid.

8.1.3 Different Types of Stock

Common Stock

When people talk about "stock," they're often referring to **common stock**. Common stock typically confers voting rights - one vote per share - though these rights can vary in **dual-** and **tri-class** share structures. Firms in which founders and families play a large role often have dual-class share structures. For example, GoPro (ticker symbol GPRO) had a dual class share structure in which "A" class shares issued to the public have one vote per share, and "B" class shares held by Nicholas Woodman, the founder, have 10 votes per share. Alphabet has a tri-class structure in which A class shares (GOOGL) issued to the public have one vote per share; B class shares, which are not traded in public markets, are held by the founders, Sergey Brin and Larry Page, and several directors of the company, and have 10 votes per share; and C class shares (GOOG) issued to the public have no voting rights. The C class shares allow Alphabet to issue more shares to raise money without diluting the voting power of the B shareholders. As of July 2023 there were 682 companies with a dual-class share structure according to the Council of Institutional Investors.

The variation in voting rights across different share classes allows a small number of investors to retain control, i.e., decision-making ability, while enabling them to raise a great deal of money from other investors. This begs the question: Why would anyone buy these shares if they're ceding control to a select few? Multiple share classes would seem to be ripe with corporate governance problems. This risk is often why we'll see these different share classes trading at different prices even though they are claims to the same cash flow streams. And, of course, shareholders can always vote with their feet by selling their shares if they don't approve of managerial decision making. In international settings, dual class shares often separate voting rights between domestic and foreign investors.

Preferred Stock

Preferred stock sits above common stock, but below debt, in the priority structure of corporate liabilities. Preferred shareholders only receive money after all debt holders have been paid, but before any common shareholders can receive money. In many ways, preferred stock is more like a bond than common stock and in the early part of the 20th century, firms relied almost exclusively on preferred stock instead of debt to finance their operations.¹ Preferred stock comes with a par value and a fixed dividend that operates like a coupon payment.

For example, as of 2023, Bank of America had several different types of preferred stock outstanding such as ticker BAC-K, non-cumulative preferred stock with a par value \$25 per share and a coupon rate of 5.875%. These terms imply an annual dividend of $25 \times 0.05875 = \$1.47$ per share, which is paid in equal quarterly installments of $1.47 / 4 = \$0.37$ per share. Non-cumulative means that if Bank of America misses a dividend payment, they are not responsible for making it up to investors. Had the shares been “cumulative” and the bank missed a dividend, they would have been required to make up the missed payment at some point and before paying any dividends to common shareholders. Unlike debt, firms cannot be forced into bankruptcy by preferred shareholders because of a missed dividend payment.

Preferred stock typically doesn't come with voting rights, though missing a dividend payment can sometimes activate voting rights for these investors. Relative to common stock, preferred stock is rare and often issued by utilities, financial institutions, and **real estate investment trusts (REITs)**, primarily for regulatory and tax purposes.

There are other types of preferred stock. **Convertible preferred** stock can be converted into common stock at the owner's discretion, and is often used by **venture capitalists** to finance startups (i.e., young companies). **Adjustable-rate preferred** stock has a dividend that adjusts each quarter according to changes in a short-term interest rate, such as the yield on a T-bill. These securities are even less common, in terms of economic scale, than traditional preferred equity.

¹See the study by Mark T. Leary, John R. Graham, and Michael R. Roberts, 2015, “A Century of Corporate Capital Structure: The leverage of Corporate America?” *Journal of Financial Economics* 118, 658-683.




Figure 8.3: Number of Listed Firms (Source: World Federation of Exchanges, March 2022)

8.2 Buying and Selling Stock

Many companies have stock, but relatively few have stock that are listed on an **exchange** - a marketplace where buyers and sellers trade financial assets.² Figure 8.3 presents data from the World Federation of Exchanges. As of March 2022, there were approximately 58,200 companies listed on exchanges worldwide. Most listed companies reside in the Asia-Pacific region - China, Japan, Hong Kong, etc. However, figure 8.4 shows the largest exchanges in terms of total **market capitalization** or **market cap** are in the U.S.. The market capitalization of a company is the value of its equity, computed as the product of the number of shares outstanding and the price per share. Market cap is *not* the value of a company, which includes the value of all claims - equity and debt.

When firms want to sell stock (or bonds) to the public, they typically work with investment banks who **underwrite** the sale. Underwriting is a process by which the underwriters - investment banks - perform several functions including the following.

- *Risk Assessment and Pricing:* Underwriters assess and analyze the risk of the securities being offered, the financial health of the company issuing the securities, the financial market conditions, and the potential demand for these securities. This assessment is used to determine a price for the securities.

²Some organizational forms, such as limited partnerships, do not have stock, only a pro rata (i.e., proportional) claim to the income of the organization.




Figure 8.4: Largest Stock Exchanges as Measured by Market Capitalization (Source: Statista, May 2023)

- *Buying and Reselling Securities:* Typically, underwriters purchase the securities from the issuer and then sell them to the public or institutional investors in the **primary market**. Buying and selling (a.k.a., trading) stock after this initial placement occurs in the secondary market on exchanges. This process involves a commitment from the underwriter to buy a specific number of shares at a predetermined price, which they then aim to sell at a profit.³
- *Guaranteeing Sale of Securities:* The nature of the guarantee depends on the type of commitment between the underwriter and the company. A **firm commitment** is where the underwriter fully commits to the offering by buying the entire issue and taking financial responsibility for any unsold shares. A **best efforts** or **marketed deal** shifts the risk of any unsold shares to the issuing company. Finally, an **all-or-none** commitment requires the underwriter to sell the entire issue at the offering price or lose their compensation for underwriting the sale.
- *Advisory Role:* Underwriters often advise the issuing company on various aspects of the securities offering, such as timing, legal compliance, and marketing strategies. This can include deciding whether to issue stocks or bonds, determining the best time to launch the offering, and advising on regulatory requirements.
- *Market Making:* After the securities are initially sold, the underwriter may also provide support in the secondary market, helping to maintain liquidity for the securities, which can be important for the issuer's reputation and future capital raising efforts.

The first time a firm sells securities to the public is called an **initial public offering** or **IPO**. Subsequent security issues are referred to as **follow-on offerings** or **seasoned equity offerings (SEOs)**.

8.2.1 Price Quotes

Like bonds, we buy and sell stock at the ask and bid prices. The difference between these prices is the bid-ask spread and corresponds to a transaction cost we pay for the ability easily buy and sell stock. Table 1 shows stock price quotes for Amazon.com Inc. (AMZN) and Iveric Bio Inc. (ISEE) as of October 21, 2022.

Size refers to the number of shares for which the price applies. For example, Amazon's quotes shows that we can sell 500 shares at \$117.25 per share or buy 100 shares at \$117.26

³A well-known exception to this step was Google's dutch auction in August 2004 in which they sold stock directly to investors.

per share. We can always buy and sell more or less shares than what's stated in the quotes. However, orders larger than the quoted sizes may be executed at different prices. For example, if we want to sell 1,000 Amazon shares, our order may be broken up and sold at different prices, assuming we don't request otherwise.

	Bid x Size	Ask x Size
Amazon.com Inc (AMZN)	117.25 x 500	117.26 x 100
Iveric Bio Inc. (ISEE)	20.94 x 200	20.97 x 400

Table 1: Amazon and Iveric Bio Price Quotes (\$/share), October 21, 2022

8.2.2 Order Types

When we want to buy or sell shares, we must specify an **order type**. There are many types of orders. We'll discuss the most common.

Market Order

A **market buy** or **market sell** order seeks the fastest execution possible, regardless of at what price the order gets executed. Take Amazon's price quotes from table 1. If we submit a market buy for 200 shares of Amazon stock, we may end up paying more or less than the \$117.26 per share if prices change before the order is executed. If we place a market buy for 1,000 Amazon shares, we could wind up purchasing shares at several different prices, some or all of which may be higher than the quoted ask price. Market sell orders face a similar risk.

Market orders minimize **execution risk** at the expense of **price risk**. A market order is likely to be executed quickly, but the price at which it will be executed is uncertain. For small orders and in normal markets, this price risk is relatively low. For volatile stocks or during turbulent markets, prices can change by large amounts very quickly so that price risk can be substantial. Likewise, large orders may themselves impact prices making a market order risky in terms of the execution price. So, market orders are most sensible when markets are normal, our order is small relative to the quoted trade sizes, and we are not concerned with the exact price at which we buy or sell.

Limit Order

A **limit order** specifies a maximum or minimum price - the **limit price** - at which shares are to be bought or sold. In doing so, limit orders minimize price risk at the expense of execution risk. For example, a **limit buy** of Iveric Bio shares at \$21 per share ensures that the order will only be executed if the shares can be purchased at a price equal to or below \$21 per share. A **limit sell** of Amazon at \$117 ensures that the order will only be executed if the shares can be sold at a price equal to or above \$117 per share.

We can choose any price for our limit order, but we should be aware that the price determines the probability of the order being executed. If we submit a limit buy of 200 shares of Amazon at \$100 per share when Amazon is currently trading at \$117.25, there is a good chance that our order will never get executed. Limit orders are useful in turbulent markets or when we are particularly sensitive to price at which our trade executes.

In addition to specifying a price for our limit order, we need to specify a term - the amount of time for which the order should be valid. Many limit orders are only valid until the end of the day on which they are submitted after which unfilled limit orders are canceled. Alternatively, we can specify a limit order be valid for a certain number of days (7, 30, 60, etc.) or indefinitely in which case the order must be canceled by the submitter or it will remain valid indefinitely. A **fill or kill** order is a limit order that must be executed in its entirety - no buying or selling a fraction of the shares requested. These are often reserved for larger trades to avoid having some portion of the order remain unexecuted.

Stop Order

Stop orders are buy and sell orders that are executed only when the price of the stock trades at or through a specified **stop price**. If the price never reaches the stop price, the order is not executed. For example, we might submit a stop buy order if we believe that a stock will continue to rise once it breaks through a certain price. For example, a stop buy of Iveric Bio at \$22 will be converted into a market buy order and executed when the price per share of Iveric Bio reaches \$22. Note that the stop price for a stop buy is *above* the current market price. If it weren't, it would be equivalent to a market buy order.

Now consider a **stop sell**, often referred to as a **stop loss**, order. We might submit a stop sell if we are concerned about a falling stock price and want to minimize losses or protect previously gotten gains. For example, a stop sell of Amazon at \$110 ensures that if the price of Amazon falls from its current level - approximately \$117 - to \$110, our stop sell order will convert to a market sell and will be executed as quickly as possible.

8.3 Making and Losing Money - Returns

Now let's understand how we we make or lose money when we trade stock, and how we commonly measure these gains and losses. Table 2 presents monthly return data from January 2021 to December 2021 for Microsoft (MSFT). The price per share column presents the **closing price** - last price at which a trade occurred - on the last day of each month on which the stock market is open. Stock markets are not open on weekends and certain holidays leaving only 252 days in a calendar year in which U.S. stock markets are open for trade.⁴ The dividend per share column presents the dividend amount that the owner of each share receives. Microsoft pays quarterly dividends.

Date	Price per Share	Dividend per Share	Price Return (%)	Dividend Yield (%)	Total Return (%)	Annual Return (%)
12/31/20	222.42					
1/29/21	231.96		4.29	0.00	4.29	
2/26/21	232.38	0.56	0.18	0.24	0.42	
3/31/21	235.77		1.46	0.00	1.46	
4/30/21	252.18		6.96	0.00	6.96	
5/28/21	249.68	0.56	-0.99	0.22	-0.77	
6/30/21	270.90		8.50	0.00	8.50	
7/30/21	284.91		5.17	0.00	5.17	
8/31/21	301.88	0.56	5.96	0.20	6.15	
9/30/21	281.92		-6.61	0.00	-6.61	
10/29/21	331.62		17.63	0.00	17.63	
11/30/21	330.59	0.62	-0.31	0.19	-0.12	
12/31/21	336.32		1.73	0.00	1.73	52.48

Table 2: Microsoft Monthly Stock Price Data

A stock return is no different from a bond return or return to any other asset. It is the ratio of the money received to money spent less one. That is,

$$\text{Return} = \frac{\text{Money we receive}}{\text{Money we spent}} - 1$$

For stocks, the money we receive is comprised of the stock price at the time of sale and any dividends. The money spent is the stock price at the time of purchase. So, the return from

⁴However, investors can trade **after hours**, outside of when the market is open.

$t - 1$ to t is

$$\text{Return}_{t-1,t} = \frac{P_t + D_t}{P_{t-1}} - 1. \quad (8.1)$$

This expression applies to both *expected* and *realized* returns. In the current context, we are looking at historical data - what's already happened. Therefore, this discussion of Microsoft is about realized returns.

Let's rewrite the equation 8.1 like so.

$$\text{Return}_{t-1,t} = \underbrace{\frac{P_t - P_{t-1}}{P_{t-1}}}_{\text{Price return}} + \underbrace{\frac{D_t}{P_{t-1}}}_{\text{Dividend yield}}$$

The **total return** on a stock is the sum of two components. The first is the **price return** (a.k.a., price appreciation/depreciation, capital gain/loss), which measures the relative change in the stock price. The second is the **dividend yield**, which measures the fraction of the current stock price paid out as a dividend. Table 2 presents monthly realizations of these two components and the total return.

For example, the monthly return for November 2021 - i.e., the end of October to the end of November - is

$$\frac{330.59 - 331.62}{331.62} + \frac{0.62}{331.62} = -0.0031 + 0.0019 = -0.0012.$$

Microsoft stock experienced a capital loss of 0.31% and a dividend yield of 0.19% for a total realized return of -0.12% during the month of November 2021. If we owned 100 shares of Microsoft stock as of October 29, our investment would have been worth $100 \times 331.62 = \$33,162$. At the end of November, our 100 shares have been worth $100 \times 330.59 = \$33,059$, but we also would have received $100 \times 0.62 = \$62$ in dividends. So, our loss on the falling share price was partially offset by the dividend payment.

8.3.1 Annual Returns

The final column shows the annual return for 2021 assuming all dividends are reinvested in Microsoft stock. In other words, any cash dividends paid to a stockholder are used to purchase more Microsoft stock. This return is computed by compounding the realized monthly returns.

$$\begin{aligned} \text{Annual Return}_{t-12,t} &= (1 + r_{t-12,t-11}) \times (1 + r_{t-11,t-10}) \times \dots \times (1 + r_{t-1,t}) - 1 \\ &= (1 + 0.0429) \times (1 + 0.0042) \times \dots \times (1 + 0.0173) \\ &= 0.525 \end{aligned}$$

For each dollar we had invested in Microsoft stock at the start of 2021, we earned \$0.525 or 52.5%. That's a hefty return!

The equation we used to get the annual return might look strange. In fact, it's just our effective annual rate (EAR) formula (equation 3.2) modified to allow the periodic rate to vary over time. Recall the definition of the EAR, r , from chapter (3).

$$r = (1 + i)^k - 1$$

The periodic interest rate is denoted by i , and the number of periods in a year by k . When we introduced the EAR, the periodic interest rate, i , was the same every period. This assumption enables us to write the EAR succinctly as $(1 + i)^k - 1$. Stock returns differ every period, so we don't get to write the expression quite as neatly even though the intuition is the same.

We can also estimate the annual return using the internal rate of return (IRR). Recall the IRR is the one discount rate such that the NPV is equal to zero, or equivalently, the current price equals the sum of the discounted cash flows. The cash flows from purchasing Microsoft stock at the end of December 2020 and selling at the end of December 2021 are illustrated in Figure 8.5.

Dates	12/20	1/21	2/21	3/21	4/21	10/21	11/21	12/21
Cash Flows	-222.42	0	0	0.56	0	0	0.62	336.32

Figure 8.5: Cash Flows from MSFT Stock Investment

We can find the monthly periodic IRR by finding IRR such that the following equality holds.

$$222.42 = \frac{0}{1 + IRR} + \frac{0}{(1 + IRR)^2} + \frac{0.56}{(1 + IRR)^3} + \dots + \frac{0.62}{(1 + IRR)^{11}} + \frac{336.32}{(1 + IRR)^{12}}$$

(These are monthly cash flows with time measured in months; hence the solution is a monthly IRR.) Using a spreadsheet program produces $IRR = 3.58\%$. The annualized IRR is

$$(1 + 0.358)^{12} - 1 = 0.5247$$

or 52.47%, which is close to the 52.48% we obtained from compounding the realized monthly returns.

The difference in these numbers is due to the IRR's assumption that all reinvested cash flows earn the same internal rate of return. In reality, Microsoft's returns varied month to

month meaning any reinvested dividends would have earned different rates of return. Thus, the IRR estimate is only an approximation of the actual return an investor received over the year. If the periodic returns are similar, the difference between the IRR and the actual return will be small. Otherwise, the difference can be quite large.

If we have a return over a period longer than a year, k is a fraction in our effective annual rate equation. For example, if we earn 164% over eight years, our annualized return is

$$(1 + i)^k - 1 = (1 + 1.64)^{1/8} - 1 = 0.0948,$$

or 9.48% per year. Note, the realized annual returns earned over the eight years that led to a 164% return may all be different from 9.48%. However, earning 164% over eight years is equivalent to earning 9.48% per year compounded annually over eight years. As with bonds or any other investment, we want to compare returns over similar horizons and annual horizons are the most common benchmark.

8.3.2 The Impact of Dividends

A dividend begins with an announcement or declaration by the company that it will pay a dividend in the near future. This usually occurs by a press release. The date of the announcement is called the **declaration date**. When declaring the dividend, the company sets a **record date** and a **payment date**.

The record date is the date on which an investor must own shares in the company in order to be entitled to the dividend. The record date also determines the **ex-dividend date** or **ex-date**, which usually occurs one business day before the record date. Investors must purchase the stock *before* the ex-date to be counted as a shareholder as of the record date and therefore entitled to the dividend. Investors purchasing on or after the ex-date are not entitled to the upcoming dividend. Finally, the dividend gets paid on the **payment date**.

Figure 8.6 illustrates the dividend payment process for Microsoft in November 2021. On September 14, Microsoft declared it would pay a dividend of \$0.62 per share with a record date of November 18. The ex-dividend date was one business day before the record date on November 17. The dividend was paid to shareholders on December 9. Investors that owned stock as of November 16 - one day before the ex-date - were entitled to this dividend.⁵ An exception occurs for especially large dividends, 25% or more of the share price. In this case,

⁵Notice in table 2 that the dividend is included in the November return despite being paid in December. The ex-date, which occurs in November, is the relevant date for determining who receives and doesn't receive the dividend.

the ex-date is one day after the payment date meaning investors can purchase the stock up to and including the day of the dividend payment and still receive the dividend.

Dates	Declaration	Ex-dividend	Record	Payment
	9/14/21	11/17/21	11/18/21	12/9/21

Figure 8.6: Dividend Dates - Microsoft (MSFT)

Finally, companies can also pay dividends in the form of shares of stock instead of cash. This stock can be additional shares of the company or of a subsidiary that is being **spun off**, i.e., sold to create a new stand alone company.

8.3.3 The Impact of Stock Splits

Occasionally, companies will replace each existing share with more than one share and adjust the price to reflect this change in the number of shares outstanding. Doing is called a **stock split** and no new money is raised or paid by the company. The rationale for stock splits is often to reduce the price of a stock that has gotten so high in the opinion of management that some investors - mostly retail investors - may not be able to purchase shares thereby depressing the stock price.⁶ For example, Berkshire Hathaway, the company owned by famed investor Warren Buffet, has class A shares that trade for over \$460,000 per share as of December 2022. Few people can afford to buy a single share of this company, though it does offer class B shares for a mere \$306. Let's see how a stock split works with an example.

Table 3 presents market data for Intuitive Surgical (ISRG). On October 5, 2021, ISRG executed a 3-for-1, sometime written as 3:1, split of their stock. Each share was converted into three shares, which can be seen in the tripling of shares outstanding from October 4 to October 5. To ensure the split has no effect on value, the price of the stock will decline by two thirds. For example, if we owned 100 shares of ISRG on October 4th, the total value of our investment would be $100 \times 970.50 = \$97,050$. Immediately after the split, we would own 300 shares but the price per share would be one third that of the pre-split price, $970.50 / 3 = \$323.50$. So, after the split the total value of our investment is $300 \times 323.50 = \$97,050$, the same as before the split. The price on October 5th is \$330.07, close but not exactly equal to \$323.50. The difference is due to trading on October 5th after the split that moves the share price off the post-split price.

⁶This rationale is less relevant today when investors can purchase **fractional shares** and therefore are no longer constrained to purchase individual shares.

Date	Price (\$/share)	Shares Outstanding (000s)	Return Unadj. (%)	Return Adj. (%)
10/4/21	970.50	118,991	-3.86	-3.86
10/5/21	330.07	356,973	-65.99	2.03
10/6/21	334.94	356,973	1.48	1.48

Table 3: Stock Split - Intuitive Surgical (ISRG)

Importantly, shareholders are no worse off following a stock split because the large price decline is exactly offset by an increase in the number of shares we own. If we naively compute the return on our investment ignoring the split, we can get some strange results as seen in the Return Unadj. column. ISRG equity did not lose almost two thirds of its value on October 5th. It split its stock 3:1. To get the correct return, Return Adj., we have to multiply the closing price by the split factor, three in the case of ISRG, like so. (There was no dividend.)

$$\text{Split Adjusted Return} = \frac{330.07 \times 3}{970.50} - 1 = 0.0203$$

Companies also execute **reverse splits** in which each share is exchanged for a fraction of a shares. Table 4 presents market data for ISRG in 2003. On July 1, the company executed a 1:2 reverse split in which every two shares that an investor owned were replaced with one share. The result was an immediate doubling of the price per share, but again no change in the wealth of investors who had half as many shares as they had before the reverse split.

Date	Price (\$/share)	Shares Outstanding (000s)	Return Unadj. (%)	Return Adj. (%)
6/30/03	7.49	37,121	-4.71	-4.71
7/1/03	14.64	18,561	95.46	-2.27
7/2/03	15.15	18,561	3.48	3.48

Table 4: Reverse Stock Split - Intuitive Surgical (ISRG)

Splits and reverse splits are common but infrequent events. Theoretically, they should have no effect on investors wealth or returns, once returns are adjusted to reflect the split. So, the doubling of ISRG's stock price on July 1, 2003 was no cause for celebration. In fact, the stock actually *lost* 2.27% that day. Likewise, the 66% decline in stock price on October 5th, 2021 is no cause for concern. In fact, shareholders earned 2.03% that day. If stock splits

have any affect on shareholder wealth it is because shareholders use splits as a signal about the future prospects of company.⁷

8.4 Historical Data and Risk-Reward Estimates

8.4.1 Expected Returns

Expected returns play a central role in valuation; they are the r in our Fundamental Value Relation. But, how can we estimate the expected return on a stock? Below and in subsequent chapters we'll discuss several approaches. The first relies on simple statistics, which tells us that the **arithmetic average**, or just **average**, can provide an estimate of what one might expect to happen in the future. In other words, we can estimate a stock's expected return looking at its average historical realized returns.

The average is computed by summing the numbers in a **sample** - collection of observations of a random phenomenon - and dividing by the number of numbers in the sample. Mathematically, the average return can be expressed as

$$\text{Average} = \frac{r_1 + r_2 + \dots + r_N}{N}, \quad (8.2)$$

where r_1, r_2, \dots, r_N are the N historical returns in our sample.

Consider estimating the expected return for Microsoft. Table 5 shows a partial view of Microsoft's historical monthly returns since its IPO.

Date	Return
19860430	0.172727
19860530	0.077519
19860630	-0.115108
⋮	⋮
20211029	0.176291
20211130	-0.001236
20211231	0.017333

Table 5: Microsoft Monthly Historical Returns

We could estimate the monthly expected return for Microsoft by using a sample of the two most recent realized returns to get $(1.73\% - 0.12\%)/2 = 0.81\%$. The problem with this

⁷See the study by David L. Ikenberry, Graeme Rankine, and Earl Stice, 2009, "What do stock splits really signal?", *Journal of Financial and Quantitative Analysis*, 31(3) 357-375.

estimate is that it is not very precise. As we'll see in a moment, stock returns vary a lot over time. If we were to estimate Microsoft's expected return using a sample of the *three* most recent returns, our estimate would be $(1.73\% - 0.12\% + 17.63\%)/3 = 6.41\%$. Adding just one data point to our sample makes a huge difference in our estimate of the average monthly return.

This reasoning might lead us to think that using as much data as possible is the best way to go. For example, we could use every monthly return since Microsoft became a publicly traded company in March of 1986 when it had its IPO. This approach gets us a sample with 431 monthly returns whose average is 2.48%. The problem with using data that's almost 40 years old is that Microsoft is a very different company today. Returns from long ago likely don't accurately reflect the risk of Microsoft's business today.

How much data should we use? Unfortunately, there is no definitive answer and the answer is likely to vary depending on application. When we estimated the returns to different investments in table 1 from chapter 1, we used data from 1927 to 2021. One might argue that because these estimates correspond to markets or securities that have been around for some time, a longer horizon is reasonable. However, the companies that comprise the market, as well as the broader macroeconomy, have changed dramatically during that time. Ultimately, there is a tradeoff between precision and relevance that is at our, the decision-maker's, discretion.

8.4.2 Volatility

If returns were all that mattered, we could just invest all of our money in the stock (or other asset) with the highest return. Of course, investing isn't that easy because we also care about risk. There are many ways to measure risk. We'll discuss several. For now, we'll focus on perhaps the most intuitive and common measure, **standard deviation** or **volatility**. Volatility measures how much returns vary around the expected return.

We saw above that we can estimate a stock's expected return by computing the average of its historical returns. We can estimate a stock's volatility with historical data by computing the standard deviation.

$$\text{Standard deviation} = \sqrt{\frac{1}{N-1} \underbrace{[(r_1 - \bar{r})^2 + (r_2 - \bar{r})^2 + \dots + (r_N - \bar{r})^2]}_{\text{Variance}}} \quad (8.3)$$

where \bar{r} is the average return and N is the number of returns. For those familiar with

statistics, the standard deviation is the positive square root of the **variance**. We focus on standard deviations because the units are easier to interpret. Variance is units squared.

The equation looks a little scary but is simple to compute. Here's a recipe, though most spreadsheet programs and programming languages have built-in functions to compute the **sample standard deviation**.

1. Compute the average of the sample returns, \bar{r} , using equation 8.2.
2. Subtract the average return from each return in the sample. This gives us **deviations from the average**: $(r_1 - \bar{r}), (r_2 - \bar{r}), \dots, (r_N - \bar{r})$.
3. Square each of these deviations from the average: $(r_1 - \bar{r})^2, (r_2 - \bar{r})^2, \dots, (r_N - \bar{r})^2$.
4. Add up all the squared deviations from the average.
5. Divide the sum by the number of returns in the sample minus one. This produces the variance.
6. Compute the positive square root of the variance.

How much data to use to estimate the volatility is again at the discretion of the decision maker. Using monthly returns for Microsoft from 2017 to 2021 produces expected return and volatility estimates of 3.11% and 5.24%. Loosely speaking, Microsoft's monthly return over this period has been on average 3.11% give or take 5.24% each month. Notice that the return volatility is larger than the average return, which tells us that the stock returns can make very large swings. This feature of stock returns is illustrated in Figure 8.7, which presents the time series of monthly returns for Microsoft over the period 2017 to 2021.

Figure 8.8 presents a histogram of Microsoft's monthly stock returns and another way of viewing just how volatile stock returns can be. Each bar corresponds to an interval of returns whose height measures the number of returns falling in that bin. For example, the left most bin corresponds to returns between -8.4% and -6.4%, of which there are four. The histogram is approximately centered over its mean of 3.11% - the dashed red line. However, the returns are spread out around that average reflecting the large volatility in returns.

To summarize, one way to estimate expected returns is to take an average of historical returns. The challenge is in identifying a sufficiently large enough sample that is also representative of the state of the company as it currently stands. This challenge is particularly difficult because returns are highly volatile; i.e., they have a large standard deviation relative




Figure 8.7: Microsoft Monthly Stock Returns




Figure 8.8: Microsoft Monthly Stock Return Histogram

to their average. Thus, estimating expected returns with historical averages produces what statisticians would call a very noisy estimate or one that could be very far from the true value. We'll see alternative approaches to estimate expected returns later in chapter 10.

8.4.3 Scaling Expected Returns and Volatility

Our expected return and volatility estimates reflect *monthly* returns because that is the frequency of the data used to estimate these quantities. What if we want to know average *annual* returns and *annual* volatility? We could just work with annual return data. Alternatively, we can multiply our monthly average and volatility estimates by appropriate scale factors.⁸ Specifically,

$$\text{Average annual return} = 12 \times \text{Average monthly return}$$

$$\text{Average annual volatility} = \sqrt{12} \times \text{Average monthly volatility}$$

The “12” and “ $\sqrt{12}$ ” in the equations above are the scale factors for converting average monthly return and volatility into average annual return and volatility. For Microsoft, the estimated annual expected return is $12 \times 3.11\% = 37.32\%$; the estimated annual volatility is $\sqrt{12} \times 5.24\% = 18.15\%$.

More generally,

$$\text{Average annual return} = k \times \text{Average periodic return} \quad (8.4)$$

$$\text{Annual volatility} = \sqrt{k} \times \text{Periodic volatility} \quad (8.5)$$

where k is the number of periods in a year. For example, when working with weekly returns, $k = 52$. When working with daily returns, $k = 252$.

⁸This scaling requires that monthly returns are **statistically independent** of one another. That is, this month's return must be unrelated to all previous (and future) monthly returns, which is true in the data.

We might wonder why we don't compound our periodic return to get an annual return like we did above. When we estimate *expected returns*, we want to use arithmetic averages (equation 8.2) and multiplying by the number of periods in a year k is sensible. When we are trying to understand *realized performance*, we want to use the compounded return. Put differently, if we want to know how our investment *did*, use the compounded return. If we want to know how our investment *will do*, use the average return, which is a statistically unbiased estimate of expected returns.

8.5 Stock Valuation

8.5.1 The Starting Point

As with all valuations, we start with our Fundamental Value Relation.

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

The cash flows correspond to the the cash flows received by shareholders. The discount rate r goes by several names in the context of stocks. **Expected stock return** denotes the return stock investors expect to earn on their investment. **Equity cost of capital** denotes the cost to companies of raising equity capital - the flip side of what equity investors expect to earn. Finally, **levered return** denotes the return that is affected by the firm's financial policy or leverage.

We can value all of a company's shares to estimate its market capitalization, in which case cash flow measures the total money flowing to *all* shareholders. Or, we can value an individual share to get the stock price, in which case cash flow measures the money flowing to a single share. Regardless of which approach we take, the two should always reconcile since the market capitalization is the price per share times the number of shares outstanding.

Unlike in bond valuation, we're going to assume that the discount rate, r , is constant over time. There are no corresponding yield curves for stocks because stocks aren't contractually required to make periodic payments to investors, and they don't have a maturity date. Additionally, the risk premium for stocks is often larger than the risk-free rate and more difficult to estimate. Recall from our discussion of credit risk in the fixed income chapter (7) that the discount rate for risky cash flows can be written like so.

$$r = \text{Risk-free rate} + \text{Risk premium}$$

To obtain estimates of the risk-free rate at different horizons, we can look at the yield curve for risk-free bonds, or something approximating risk-free bonds like Treasury securities in the U.S. The risk premium must be estimated using historical data or forecasts, and as we'll see in subsequent chapters, this is challenging. So, we follow practice and assume the expected stock return is constant, but possibly updated with the passage of time.

8.5.2 The Effect of Investment Horizon on Value

To set the stage for the remainder of this chapter, let's explore an interesting implication of our Fundamental Value Relation, namely, that the value of an asset is the same for short-term and long-term investors. Put differently, how long we intend to invest in an asset doesn't matter for its value to us.

This implication flies in the face of conventional wisdom and a lot of popular press that believes short-term investors destroy value. The argument is that short-term investors are only interested in extracting financial value for themselves, and they do so at the expense of long-term investors. **Private equity investors** and **activist investors** are often accused of such short-termism and value-destruction. Both of these investors acquire stock to enact changes in governance, operations, strategy, and financing at companies for the ultimate goal of selling their stock at a higher price. Their investment horizons can vary from a few months over ten years, but both seek to exit their investments as soon as possible.

With a short investment horizon, it would seem these investors would do everything to increase cash flows in the near-term, including sacrificing cash flows in the long-term. More precisely, short-term investors are often accused of increasing $CashFlow_1, \dots, CashFlow_T$ by sacrificing all cash flows beyond their investment horizon, $(CashFlow_{T+1}, CashFlow_{T+2}, \dots)$ But, does this argument hold water? Is it rational for short-term investors to take actions that reduce the cash flows, and therefore value, of the company after their investment horizon?

Consider selling a stock one period (day, month, year, etc.) after purchase. According to our Fundamental Value Relation, the value of the stock ($Value_0$) should equal the sum of the discounted future cash flows consisting of any dividend the stock may pay (Div_1) and the sale price of the stock ($Price_1$).

$$Value_0 = \frac{Div_1 + Price_1}{(1+r)} \quad (8.6)$$

To be clear, the dividend, sales price, and stock return (r) are all *expected* values. We don't know what they are today. They occur in the future.

In equation 8.6, we're assuming that the dividend, if any, is paid just before we sell the stock. In finance lingo, we are selling the stock **ex-dividend**, i.e., without the dividend. Had we sold the stock just before receiving the dividend, we would say the stock is **cum-dividend**, i.e., with the dividend. This assumption is made to avoid the extra calculation of having to discount the dividend and the price by different numbers of time periods. If the stock doesn't pay a dividend before we sell it, then $Div_1 = 0$.

The question now is: At what price can we sell our stock one period from today, $Price_1$? The person to whom we sell the stock will pay us the present value of *their* future cash flows. Assuming they sell their stock after one period, the price they should be willing to pay is as follows.

$$Value_1 = Price_1 = \frac{Div_2 + Price_2}{1 + r}$$

Substituting this expression for $Price_1$ in equation 8.6 produces the following relation for the value of the stock today.

$$Value_0 = \frac{Div_1 + Price_1}{(1 + r)} = \frac{Div_1 + \overbrace{\left(\frac{Div_2 + Price_2}{1 + r}\right)}^{Price_1}}{(1 + r)} = \frac{Div_1}{1 + r} + \frac{Div_2 + Price_2}{(1 + r)^2}$$

The price in period 2 should equal the present value of the future cash flows or

$$Value_2 = Price_2 = \frac{Div_3 + Price_3}{1 + r}.$$

Substituting again produces

$$Value_0 = \frac{Div_1}{1 + r} + \frac{Div_2 + \overbrace{\left(\frac{Div_3 + Price_3}{1 + r}\right)}^{Price_2}}{(1 + r)^2} = \frac{Div_1}{1 + r} + \frac{Div_2}{(1 + r)^2} + \frac{Div_3 + Price_3}{(1 + r)^3}.$$

Are you seeing a pattern?

If we keep replacing the future price with its future cash flows for an arbitrary number of periods, T , we get that the value of the stock today is

$$Value_0 = \frac{Div_1}{1 + r} + \frac{Div_2}{(1 + r)^2} + \frac{Div_3}{(1 + r)^3} + \dots + \frac{Div_T + Price_T}{(1 + r)^T}. \quad (8.7)$$

If we keep replacing the price forever, we get our first model of stock valuation, the **dividend discount model** or **DDM**. The DDM says that the value of a stock equals the sum of the discounted future dividends in perpetuity.

$$Value_0 = \frac{Div_1}{(1 + r)} + \frac{Div_2}{(1 + r)^2} + \frac{Div_3}{(1 + r)^3} + \dots \quad (8.8)$$

While no company exists forever, our analysis shows that we can treat a company as if it does. For example, if a company is acquired, the price of the acquisition is based on the future cash flows. Likewise, if a company goes bankrupt, its future cash flows go to zero.

Returning to the motivation for this analysis, we've just shown that it doesn't matter for value if investors are short-term, long-term, or a mix. The value of the stock at any point in time is *always* the sum of the discounted future cash flows, dividends in this case. To whomever we sell the stock, they will value the stock by looking at their future discounted cash flows. Similarly, when that person sells the stock, whomever buys it will value the stock by looking at their future discounted cash flows. And so on. Consequently, short-term investors are no different from long-term investors. If they want to sell their stock for a high price, the future cash flows *after* the sale matter because they affect the sale price.

Reconsider the conventional wisdom mentioned above. It simply doesn't make sense that activist or private equity investors sacrifice long-term cash flows in the pursuit of short-term gains. These investors make money in large part by selling their shares at a higher price in the future. Because financial markets are so competitive, there are relatively few investors that would pay a higher price for a company that has worse prospects (i.e., lower cash flows) than it did before the activist or private equity investor came on board.

This isn't to say that activist investors and private equity investors don't want to make money over short-horizons. They do and part of their strategy is often increasing dividends in the short-run. However, if they want to successfully exit their investment by selling at a higher price than at what they bought, they must create value for future investors by making changes to the firm that improves long-run cash flow prospects or hope that future investors will be fooled into buying a lemon. Likewise activist and private equity investors can make decisions that sacrifice long-term cash flows. But this is more often a result of bad decision-making than intentional value destruction. Ultimately, the business model of these investors can't be the destruction of "long-term value" in the pursuit of short-run gains because, as we've shown, the two are intimately related.

8.6 Valuation: The Dividend Discount Model (DDM)

The dividend discount model in equation 8.8 above assumes that the cash flows in our Fundamental Value Relation consist of future dividends and the discount rate is the expected return on the stock. But, what about stocks that don't pay dividends? Between 1990 and 2022, the fraction of publicly listed U.S. firms paying a dividend has never gotten above 35%

and averaged 31%. This result begs the question: If a company doesn't pay a dividend, how can it have value? Or, because most firms don't pay dividends, is the DDM practically relevant?

In fact, the DDM is relevant and insightful, albeit in a limited manner. Remember, what drives value in this model is *all* of the future dividends, not just the current dividend (or next few dividends). Even if a company doesn't pay a dividend for a long time, that doesn't mean it is worthless, only that those dividends are far off into the future. It took Microsoft 17 years after they became a public company to pay their first dividend. **Liquidating dividends** are often paid to shareholders when a company is dissolved, so even for companies that struggle, money is often returned to shareholders.

However, unless we have a lot of time, the dividend discount model is impossible to implement without additional assumptions because we need to forecast dividends *forever*. The remainder of this section explores some common assumptions to get around this problem and their implications for stock valuation.

8.6.1 Valuing Microsoft with the Gordon Growth Model

The most common assumption is that dividends grow at a constant rate, g . In this case, the dividend discount model becomes a growing perpetuity and is referred to as the **Constant Dividend Growth Model** or the **Gordon Growth Model** after Myron Gordon.

$$Value_0 = \frac{D_1}{(1+r)} + \frac{D_1(1+g)}{(1+r)^2} + \frac{D_2(1+g)^2}{(1+r)^3} + \dots = \frac{D_1}{r-g} \quad (8.9)$$

This model requires dividends grow at a constant rate, g , *forever*. As such, g is referred to as the **sustainable growth rate**. The model also requires the growth rate be strictly less than the discount rate r .⁹ As we saw in figure 8.1, dividends don't grow at a constant rate. They can even be quite erratic. Nonetheless, this model is often a useful starting point for valuations and can provide several insights.

To use the constant growth model for valuation, we need to know three things:

1. dividend growth rate, g .
2. expected stock return, r , and

⁹If $g > r$, then dividends are growing more quickly than the rate at which they're discounted and value is infinite.

3. *next* period's dividend, Div_1 .

Because all of these quantities are realized in the future, we have to estimate them. Alternatively, we can use the the current stock price, if one exists, to back out the implied growth rate or expected return. There are many ways to estimate each input and many different data sources for helping us do so. Below we highlight several approaches using publicly available data from Yahoo! Finance. Strengths and weaknesses of each estimate are noted.

Utilizing different estimation strategies and data sources isn't simply busy work. It's a part of sensitivity analysis, which we highlighted in the context of capital budgeting (chapter 5). Every estimate we produce will be wrong. The questions are: How wrong, and why is it wrong? The answers to these questions are the real value of having a model because it forces us to understand and clearly articulate a rationale for the valuation it produces and, ultimately, any decision we make.

As our illustrative vehicle, we'll use Microsoft. On July 14, 2023, its stock price was \$311.21 per share and its most recent accounting information is from March 31, 2023.

Estimating dividend growth, g .

- *Average historical growth.* Microsoft's average annual dividend growth between January 2018 and December 2022 has been 10.27%. The choice of sample, 2018 to 2022, attempts to balance a large enough sample to get a precise estimate with a relevant sample that reflects the company as it is today and, hopefully, will be in the future. This estimate is simple and empirically based, albeit on a somewhat arbitrarily chosen set of data. Additionally, the past may not accurately represent the future.
- *Analysts forecasts.* Both buy-side and sell-side analysts monitor Microsoft for their companies and clients.¹⁰ Part of their job is making forecasts about Microsoft's earnings over the next several years. The average of their earnings forecast suggests Microsoft's earnings will grow by 12.79% per year over the next five years. Assuming Microsoft distributes a constant fraction of its earnings in dividends, then it's dividends will grow by the same rate.¹¹ The fraction of earnings (a.k.a., net income)

¹⁰Buy-side analysts work for firms that purchase securities, such as hedge funds and institutional investors. Sell side analysts work for firms that sell securities, such as investment banks.

¹¹Imagine earnings, $Earn$, grow at rate g each year. Then annual earnings are $Earn_1, Earn_1(1+g), Earn_1(1+g)^2, \dots$. If we distribute a constant fraction, d , of earnings each year, then the dividend stream is $dEarn_1, dEarn_1(1+g), dEarn_1(1+g)^2, \dots$. In other words, dividends grow at the same rate, g , as earnings.

distributed in dividends is the **dividend payout ratio** or more simply **payout ratio**, which we'll denote by d .

$$\text{Payout ratio} = d = \frac{\text{Dividends}}{\text{Earnings}} \quad (8.10)$$

Microsoft's dividends and earnings were \$18.964 and \$67.45 billion, respectively, implying a payout ratio of $18.96 / 67.45 = 28.12\%$. This estimate has the advantage of incorporating analyst research aimed at capturing future trends and changes to the company, industry, and broader economy. It therefore doesn't rely entirely upon the future looking like the past as with a historical average. However, it incorporates any biases analysts bring to these forecasts.

- *ROE and retention rate.* Microsoft can do two things with its earnings: (1) distribute them to shareholders, and (2) reinvest them. It's the latter that generates growth as long as Microsoft invests these earnings in positive NPV projects or, loosely speaking, projects with a rate of return greater than their cost of capital, r .

The return on equity (ROE) measures how much earnings grow for each dollar of invested equity. Microsoft's ROE (earnings divided by book equity) was \$67.45 billion divided by \$194.68 billion = 34.65%. Each dollar of invested earnings grew by \$0.3456. However, Microsoft wasn't reinvesting every dollar of earnings; they were distributing 28.12% in dividends implying a **plowback ratio** or **retention rate** of

$$\text{Plowback ratio} = (1 - \text{Payout ratio}) = (1 - d) = 1 - 0.2812 = 0.7188. \quad (8.11)$$

Therefore, earnings are expected to grow by $0.7188 \times 0.3456 = 24.90\%$. If Microsoft maintains a constant payout ratio, dividends will grow at the same rate.

More generally, earnings growth can be estimated with the following expression.

$$\text{Earnings growth} = ROE \times \underbrace{(1 - d)}_{\text{Retention rate}} \quad (8.12)$$

(See the technical appendix for a derivation.)

This estimate has the advantage of linking the dividend growth to the investment policy of the firm. However, it does so in a rather naive manner and relies on limited information. For example, if the most recent return on equity or retention rate is temporarily high or low, the growth estimate could be biased and significantly so. Taking averages of recent ROE and retention rates can help mitigate this problem but suffer from the criticism that the future need not look like the past.

Estimating expected equity return, r .

- *Average historical returns.* Microsoft's average annual stock return from January 2018 to December 2022 was 31.19%.¹² This estimate is simple to compute and easy to understand, but relies on the future being similar to the past and is very noisy or imprecise.
- *Asset pricing model.* We could use a formal model of returns, such as the **Capital Asset Pricing Model (CAPM)** or a more general **multifactor model**. We'll discuss these models in detail later in chapter 10. For now, we'll simply use the estimate generated from this model, 8.47%, as an alternative point of comparison. Model-based estimates such as the CAPM have the advantage of being grounded in financial theory but the disadvantage of being reliant on the simplifying assumptions behind of the models. They also rely on historical data.

Estimating next period's dividend, Div_1 .

- *Historical average.* The historical average over the 2018 to 2022 period \$2.08 per share. However, this is a poor estimate because Microsoft's dividend has been trending up over time so historical dividends levels are clearly not reflective of future dividend levels. In general, using historical averages to forecast variables that trend up or down over time - what statisticians call **nonstationary** time series - is a bad idea.
- *Applying a growth estimate to current dividend.* Next period's dividend can be estimated by applying an estimate of dividend growth to the current dividend of \$2.55 per share.

$$Div_1 = Div_0 \times (1 + g)$$

If we use historical average growth, our estimate of next period's dividend is $2.55 \times (1 + 0.1027) = \2.81 per share. Using analyst forecasted growth produces a slightly larger estimate, $2.55 \times (1 + 0.1279) = \2.87 .

- *Analysts forecast.* The average analyst forecasts a **forward dividend rate**, i.e., future dividend per share, of \$2.72. As with earnings growth, analyst forecasts incorporate assessments of future performance, as well as any biases.

¹²Monthly rolling annual returns were averaged over this period.

Table 6 summarizes these estimates and highlights just how much they can differ. Table 7 shows the implications of these different estimates by computing the implied share prices for all 18 different combinations of these inputs using the Gordon Growth model.

$$Value_0 = \frac{Div_1}{r - g}$$

	Historical Average	Analyst Forecast	Model-based Estimate
Dividend growth, g (%)	10.27	12.79	24.90
Expected return, r (%)	31.19		8.47
Next period dividend, D_1 (\$/share)	2.08	2.81	2.72

Table 6: Constant Dividend Growth Model Input Estimates for Microsoft Corp. as of July 2023

g (%)	10.27	10.27	12.79	12.79	24.90	24.90
r (%)	31.19	8.47	31.19	8.47	31.19	8.47
D_1 (\$/share)						
2.08	9.94	N/A	11.30	N/A	33.09	N/A
2.81	13.44	N/A	15.28	N/A	44.73	N/A
2.72	13.00	N/A	14.78	N/A	43.29	N/A

Table 7: Microsoft Corp. Stock Price Estimates from the Constant Dividend Growth Model as of July 2023

The valuations warrant some comments. First, half of the valuations are “N/A,” or not applicable, because the growth rate is larger than the discount rate. In these cases, the growing perpetuity formula (equation 2.7), and by extension the Gordon Growth Model, are invalid and the value of the asset is infinite. Since no asset is worth an infinite amount of money, these combinations of inputs can be dismissed.

The other estimates are all positive and range from \$9.94 to \$44.73 per share. This is a wide range, but still a very long way from the current market price of \$311.21. Why? There are several possible reasons.

1. We underestimated next year’s dividend, Div_1 . But, assuming the most aggressive dividend growth of 24.90% and an expected return of 31.19%, next year’s dividend per

Time Period	Nominal	Real
1948-1979	7.71	3.87
1980-2022	5.48	2.57
1948-2022	6.43	3.12

Table 8: U.S. Real and Nominal Average GDP Growth. Source: FRED

share would have to be almost \$20 to justify a current share price of \$311.21.

$$Value_0 = \frac{Div_1}{r - g} \implies Div_1 = Value_0 \times (r - g) = 311.21 \times (0.3119 - 0.2490) = \$19.55.$$

This is an extraordinarily large dividend. To put this amount in clearer perspective, a \$19.55 dividend per share corresponds to a distribution of over \$145 billion to shareholders - more than twice their most recent earnings of \$67 billion.¹³ However, even if Microsoft decided to pay a one-time, **special dividend** in this amount, that alone is not enough to justify the current price. Microsoft would have to grow that \$19.55 per share dividend by 24.9% per year *forever*. In just 10 years, they would be paying more than \$180 per share in dividends, or \$1.3 trillion to their shareholders. This is unlikely.

2. We underestimated the dividend growth rate, g . However, our estimates of dividend growth - 10.27% and 24.90% - are in fact far too *high* to be plausible. The growth rate in the Gordon Growth Model describes how dividends grow *forever*. But, as Table 8 shows, the U.S. economy has historically grown much more slowly. The average nominal growth rate over the last 74 years has been 6.43%. More recently, post-industrial revolution, nominal growth has slowed to 5.48%.

To sustain even a 10% growth rate in perpetuity would mean that Microsoft would quickly become larger than the rest of the entire U.S. economy and the global economy shortly thereafter. Clearly, this is impossible and the data in table 8 provide a useful cap on the perpetual growth of *any* company. No company's nominal cash flow growth should outpace the nominal growth of the economy in which it operates. Thus, 6.43%, and 5.48% are upper limits for perpetual cash flow growth in any model. The same logic applies to real growth if we are working with real cash flows.

3. We overestimated the expected equity return, r . The estimate of 31.19% in perpetuity is almost surely too high. Microsoft has performed exceedingly well over recent years as

¹³There were 7.44 billion shares outstanding; hence $19.55 \times 7.44 = \$67$ billion.

a result of growing demand for their cloud services and the shift to work-from-home as a result of the Covid crisis in 2020. This is in stark contrast to its average return from June 2001 to June 2005, which was -6.11%. However, even with a more reasonable, model-based estimate of 8.47%, there is no way to achieve a price of \$311.21 given the forecasted dividend and dividend growth assumptions.

4. More than one of our assumptions is incorrect. For example, consider a dividend growth rate equal to long-term nominal economic growth of approximately 6.17%, and an expected return of 7%. This combination of assumptions, in conjunction with an analyst forecasted dividend of \$2.81 per share, produces a estimate closer to the current market price of \$311.21.

$$Value_0 = \frac{2.81}{0.07 - 0.0617} = \$338.55$$

However, the dividend growth assumption represents a sharp and unlikely departure from recent trends, especially given Microsoft's success in cloud services and Artificial Intelligence - two spaces that are experiencing tremendous growth in 2023 and 2024. Similarly, yields on Treasury securities - proxies for risk-free investments - are between 4% and 5.5%, meaning an expected return of 7% on Microsoft stock represents a very low risk premium. This is all to say that while our assumptions were able to generate a price close to what we see in the market, they appear to be at odds with reality.

5. Perhaps our estimates are correct, and the market is wrong. Perhaps Microsoft is **mispriced**, overvalued to be precise. Put differently, Microsoft's current stock price is not justified by their fundamentals - cash flows, growth, and discount rates.
6. The model is wrong. The Gordon Growth model makes three assumptions about stock prices that impose restrictions on our Fundamental Value Relation.
 - (a) Cash flows to shareholders are measured by dividends.
 - (b) Dividends grow at a constant rate, g , forever
 - (c) Expected returns, r , are constant over time.

If one or more of these assumptions is incorrect, the model will have difficulty accurately estimating value. The degree of difficulty will depend on how "wrong" the assumption(s) is. We mentioned that Microsoft is growing at a high rate, which cannot possibly be sustained forever, suggesting that the second assumption problematic. Additionally, as we'll see below, Microsoft distributes a large amount of money to shareholders in the form of share repurchases as opposed to dividends. Thus, the first

	Historical Average	Analyst Forecast	Model-based Estimate
Dividend growth, g (%)	0.26	6.30	2.87
Expected return, r (%)	16.31		6.51
Next period dividend, D_1 (\$/share)	1.46	1.36	1.44

Table 9: Constant Dividend Growth Model Input Estimates for Exelon Corp. as of July 2023

	g (%)	0.26	0.26	6.30	6.30	2.87	2.87
	r (%)	16.31	6.51	16.31	6.51	16.31	6.51
D_1 (\$/share)							
	1.46	9.11	23.38	14.60	684.98	10.88	40.16
	1.36	8.44	21.66	13.53	634.78	10.08	37.22
	1.44	8.97	23.02	14.38	674.57	10.71	39.55

Table 10: Exelon Corp. Stock Price Estimates from the Constant Dividend Growth Model as of July 2023

assumption is problematic as well. Before turning to alternative models, let's first examine a case in which the constant dividend growth model does a better job.

8.6.2 Valuing Exelon with the Gordon Growth Model

Let's repeat the valuation exercise this time using Exelon Corp. (EXC), a gas and electric utility holding company covering parts of the mid-Atlantic region of the United States (Pennsylvania, Maryland, Virginia, New Jersey). As of July 15, 2023, Exelon's stock price was \$42.02. Table 9 contains Exelon Corp.'s estimated inputs for the Gordon Growth model. The estimates are produced in a similar fashion to those for Microsoft. Table 10 contains the corresponding valuations.

$$Value_0 = \frac{Div_1}{r - g}$$

Several comments are in order. There is no scenario in which an estimated discount rate is less than an estimated dividend growth rate. Forecasts of next period's dividend, Div_1 , are all similar and not a the primary source of variation in the valuations. The scenario in which $r = 6.51\%$ and $g = 6.30\%$ leads to extremely large valuations - 15 times larger




Figure 8.9: Exelon Corp. Stock Price as a Function of the Expected Return less the Dividend Growth Rate

than the current stock price - because the difference between the two estimates is so small. We saw this phenomenon above with Microsoft. As the gap between the expected return and growth rate, $r - g$, shrinks, valuations increase nonlinearly; we're dividing by something approaching zero. Economically, a small difference between r and g is saying that dividends are growing at a rate similar to the rate at which investors are discounting them. Figure 8.9 illustrates this relation.

The historical average growth rate is not a good proxy for future growth in this case because of Exelon's erratic dividend trajectory illustrated in Figure 8.10. Exelon cut their quarterly dividend from \$0.42 to \$0.345 per share in February 2018 after which it gradually began increasing its dividend until February of 2022 when it cut its dividend again from \$0.3825 to \$0.3375 per share. So the question is whether the analyst or model-based growth rate forecast is more appropriate, and arguments can be made for and against both. The analyst forecast is close to the long-run nominal economic growth implying Exelon will grow with the economy as it serves the demand for energy. The model based estimate suggests it will gradually get smaller as the economy grows, perhaps because of the increasing importance of renewable energy which may not play a large enough role in Exelon's long-term strategy. Time will tell.

Finally, Exelon's expected equity return is unlikely its historical average of 16.51%. Like Microsoft, this high average in large part represents a favorable economic environment for their business. Being a utility, Exelon is relatively insensitive to changing economic condi-




Figure 8.10: Exelon Corp. Quarterly and Annual Dividend per Share

tions and therefore a relatively low risk investment. The model-based estimate of 6.51% is more plausible and, when coupled with the model-based growth rate, produces price estimates that are close but below the current market price. With the exception of when r and g are almost equal, all of the estimates are below the current market price suggesting that Exelon may be overvalued by the market.

In sum, the Gordon Growth Model appears to do a better job estimating equity values for Exelon than Microsoft. We'll explore why as we continue this chapter. Before then, it's useful to recognize the following anecdote about a financial consultant hired to perform a corporate valuation. The client asks the consultant: "What's the value of my company?" To which, the consultant replies: "What do you want the value of your company to be?" We can get any valuation we want with the right assumptions about the model inputs. This concern applies not only to the dividend discount model but *any* financial model in any setting - capital budgeting, bond valuation, corporate valuation, etc. So, how can we ensure we are not just making numbers up to achieve an outcome we or others desire?

The reliability of any valuation comes down to the appropriateness of the model and defensibility of its assumptions. Can we stand in front of a room of experts and make a compelling case for our model and assumptions - and therefore the estimated value? If we can't, we should reconsider using its results.

8.6.3 Learning from Prices

What if we flip the script? Instead of using the Gordon Growth model to value Microsoft and Exelon, let's use their stock prices and the model to infer different value drivers, such as the dividend growth rate and expected return, to determine whether they are economically plausible. This is often a useful first step in determining whether current market prices reflect intrinsic value.

Price Implied Expected Return

We can start with the implied expected return.

$$Price_0 = \frac{Div_1}{r - g} \implies r = \frac{Div_1}{Price_0} + g \quad (8.13)$$

Equation 8.13 says the expected return on a stock equals the expected (a.k.a., **forward** or **prospective**) dividend yield, $Div_1/Price_0$, plus the dividend growth rate, g .

As a brief aside, it's worth noting that the expected stock return can also be derived as follows assuming the current price accurately reflects the value of the stock.

$$Price_0 = \frac{Div_1 + Price_1}{1 + r} \implies r = \frac{Div_1}{Price_0} + \frac{Price_1 - Price_0}{Price_0} \quad (8.14)$$

Equation 8.14 says that the expected return is the sum of the expected dividend yield and price appreciation or capital gains. It is analogous to equation 8.1 which expressed *realized* returns in terms of the realized dividend yield and relative price change. Setting equations 8.13 and 8.14 equal implies that price growth comes from dividend growth. This is why (profitable) growth is so important to shareholders; it leads to increases in their stock price.

Let's use the model-based estimates for next-year's dividends and analysts's forecasted growth rates to back out the expected equity returns for Microsoft and Exelon.

$$\begin{aligned} \text{Microsoft: } r &= \frac{Div_1}{Price_0} + g = \frac{2.72}{311.21} + 0.1279 = 0.1366 \\ \text{Exelon: } r &= \frac{Div_1}{Price_0} + g = \frac{1.44}{42.02} + 0.0630 = 0.0973 \end{aligned}$$

Microsoft's estimate is far more reasonable than the historical average of 31.19%, but still quite a bit higher than the model-implied estimate of 8.47%. Exelon's implied expected return falls between its average historical return (16.31%) and model-implied estimate (6.51%).

Are these economically plausible estimates? We need a reference point and several were given in table 1 in chapter 1. The average annual return on the 10-year Treasury and the

S&P500 are 5.11% and 11.82%, respectively. In other words, a long-term risk-free investment has returned 5.11% per year, while the stock market as a whole has produced 11.82%. Where do Microsoft and Exelon fall on this spectrum of risk?

We'll answer this question formally later, but for now let's use some financial intuition. Clearly, both stocks are riskier investments than a Treasury security, which is guaranteed to pay a fix stream of cash flows. So, both stock should have expected returns higher than that of a Treasury. Are they riskier than the market? Possibly Microsoft, though they are so large and ubiquitous in business with their software one could argue that investing in Microsoft is much like investing in the broader economy.

Exelon, however, is arguably less risky than investing in the broad stock market. As a utility, their cash flows are relatively stable and predictable because they are governed in large part by demographics and weather. Additionally, as a natural monopoly, Exelon is regulated by both state and federal government. While climate change is reducing weather predictability and regulatory uncertainty exists, utility cash flows are arguably less risky than those generated by the broader market.

In sum, both implied expected returns are defensible *but* to get these estimates we had to assume a growth rate. While Exelon's growth rate of 6.3% is not far out of line with long-run nominal economic growth, Microsoft's 12.79% is more than double it. This is simply not possible and cause for suspicion that either the price is wrong, the other estimates are wrong, or the model itself is wrong.

Price Implied Dividend Growth

If we instead start with an assumed expected return and forecasted dividend, we can use the Gordon growth model to back out the expected dividend growth like so.

$$Price_0 = \frac{Div_1}{r - g} \implies g = r - \frac{Div_1}{Price_0} \quad (8.15)$$

Let's assume that Microsoft and Exelon's expected returns are 8.47% and 6.51%, and that their forecasted dividends per share are \$2.72 and \$1.44. (See tables 6 and 9 for these estimates.) The implied dividend growth rates are

$$\text{Microsoft: } g = r - \frac{Div_1}{Price_0} = 0.0847 - \frac{2.81}{311.21} = 0.0757, \text{ and}$$

$$\text{Exelon: } g = r - \frac{Div_1}{Price_0} = 0.0651 - \frac{1.44}{42.04} = 0.0309.$$

Microsoft's implied growth rate is still high relative to long-run nominal economic growth, while Exelon's is somewhat low. The latter result is plausible, suggesting that Exelon is growing slightly slower than the broader economy. A near 8% sustainable growth rate; however, is unsustainable and implausible for companies.

Price Implied Forward Dividend

We could also back out next period's dividend by taking a stance on expected returns and dividend growth. However, doing so is somewhat less interesting. Dividends are relatively easy to predict. Any shortcomings of our analysis thus far is unlikely a result of wildly inaccurate dividend forecasts. The challenge in applying the Gordon Growth model - all financial models - is in estimating cash flow growth and expected returns.

8.6.4 Value Decomposition: Assets in Place vs. Growth Options

What fraction of Microsoft's and Exelon's value comes from its assets in place versus its growth opportunities? Recognizing that the forward dividend, Div_1 , can be written as the product of the payout ratio, d , and **forward earnings**, $Earn_1$, we can write the Gordon Growth Model as follows as long as the company maintains a constant payout ratio.

$$Price_0 = \frac{d \times Earn_1}{r - g} \quad (8.16)$$

Now let's add and subtract $Earn_1/r$ to the right side.

$$Price_0 = \underbrace{\frac{Earn_1}{r}}_{\text{No-growth value}} + \left(\underbrace{\frac{d \times Earn_1}{r - g} - \frac{Earn_1}{r}}_{\text{Value from growth}} \right) \quad (8.17)$$

Equation 8.17 decomposes the price of a stock into two components. The first component is the present value of the assets of the company assuming no additional investments are made beyond maintaining the assets. If the company just maintains its assets, it won't grow, and its value is simply the present value of a non-growing perpetuity - what we've labeled as "No-growth value." The second component is the value from investing beyond maintenance of existing assets - Value from growth. The second term is referred to as the **present value of growth opportunities** or **PVGO**.

Assuming a forward **earnings per share EPS** of \$10.22 and equity cost of capital of 8.47%, Microsoft's stock price can be decomposed like so.

$$311.21 = \frac{10.22}{0.0847} + PVGO = 120.73 + 190.48$$

A similar calculation for Exelon produces

$$42.02 = \frac{2.21}{0.0651} + PVGO = 33.96 + 8.06.$$

What these decompositions show is that the majority ($190.48/311.21 = 61\%$) of the market's assessment of Microsoft's value comes from its future investments and growth potential. In contrast, only $8.06/42.02 = 19\%$ of Exelon's current market value comes from its future investments and growth potential. These are not unreasonable estimates. Microsoft's upside potential with cloud-based computing and its recent advances in artificial intelligence (AI) suggest strong growth potential, despite its enormous scale - over \$2.0 trillion market cap as of July 2023. In contrast, Exelon's revenue comes almost entirely from fossil fuels, which do not represent a great opportunity for growth as of 2023.

To summarize, the Gordon Growth Model is a simple but useful model for stock valuation and for understanding the value drivers behind any valuation. However, its simplicity is both a strength and weakness. As we saw with Microsoft, the model struggles to generate a valuation remotely close to the current market valuation. Similarly, the price-implied model parameters - growth rate, expected return, and dividend, are implausible. In contrast, the Gordon Growth Model did a reasonable, if not perfect, job of describing the valuation of Exelon. Why the difference? Microsoft, unlike Exelon, is growing too quickly as of 2023 to assume a constant perpetual growth as the Gordon Growth Model does.

8.6.5 Price-to-Earnings (P/E) Ratio

Equation 8.16 showed that the Gordon Growth Model can be written in terms of earnings, instead of dividends, if we assume the firm pays out a constant fraction, d , of earnings in dividends.

$$Price_0 = \frac{d \times Earn_1}{r - g}$$

Dividing both sides of this equation by earnings, $Earn_1$, generates an expression for the **forward or prospective price-to-earnings ratio** (i.e., **P/E ratio**).

$$\frac{P_0}{Earn_1} = \frac{d}{r - g} \tag{8.18}$$

The P/E ratio measures how much investors are willing to pay for \$1 of next year's expected earnings and is an example of a **valuation multiple** or **market multiple**.

Equation 8.18 shows that the P/E ratio is a function of three features of a firm.

1. The payout ratio, d . As the payout ratio increases, the P/E ratio increases, all else equal. Investors like receiving more cash flow and therefore pay more today.
2. The cost of equity capital (r). As the expected equity return increases, the P/E ratio decreases, all else equal. Investors don't like higher risk and therefore pay less today.
3. The earnings growth rate, g . As earnings grow more rapidly, the P/E ratio increases, all else equal. Investors like earnings growth, which leads to higher future prices and therefore pay more today.

The "all else equal" caveat is important because in practice one cannot expect to change one feature without affecting another. For example, simply increasing the payout ratio will almost surely negatively affect earnings growth because the firm is distributing more and investing less money.

There are several ways to measure the P/E ratio. All require use of the market capitalization or stock price in the numerator. For the denominator, we could use the total earnings or EPS over the most recent year (i.e., **trailing twelve months** or **TTM**). However, if next year's earnings are expected to differ significantly, using historical earnings can be misleading. Ideally, we would use a forecast of next year's earnings or EPS.

Let's examine Microsoft's trailing twelve month P/E ratio, in the context of its peers, the broader stock market, and its history. Table 11 presents trailing twelve month P/E ratios for 2019 - 2022. The 2022 figures are measured as of February 10, 2022. The other figures are measured as of March of each year, except Oracle which is measured as of February because of its May fiscal year end.

Microsoft's P/E ratio has been increasing over the last four years, indicating that the market is willing to pay more for a dollar of earnings. According to equation 8.18, this increase is driven by a larger payout ratio, greater earnings growth, lower expected return or some combination. Precisely which requires further analysis of the Microsoft's business model, its competitive environment, and the broader economy.

We see even more variation in P/E ratios across companies. For example, in 2020, P/E ratios ranged from a low of 9.5 for IBM to a high of 93.2 for Amazon. Does this mean

	2019	2020	2021	2022
Microsoft	25.4	25.9	31.9	32.3
Market	21.0	25.0	34.0	26.3
Amazon	74.3	93.2	58.8	59.9
Google	29.5	23.5	27.5	25.9
Apple	15.5	19.6	27.2	30.1
Oracle	18.1	15.3	15.2	23.3
IBM	12.3	9.5	20.3	25.1

Table 11: Trailing Twelve Month P/E Ratios

that Amazon is a “better” company than IBM? Not necessarily, though Amazon’s recent performance has been much better than IBM’s. Equation 8.18 tells us that the market expects greater payouts and earnings growth, or lower expected returns from Amazon relative to IBM.

Does this mean we should short Amazon (low expected return) and buy IBM (high expected return)? More generally, are stocks with high P/E ratios necessarily overpriced and stocks with low P/E ratios underpriced? No. Different firms may simply have different growth and risk expectations - g and r - as a result of different business models and different financial policies (i.e., the mix of debt and equity that companies use to fund their operations).

The companies in Table 11 all have somewhat similar capital structures consisting of primarily equity financing. So, the question remains: Should we buy the low P/E firms and (short) sell the high P/E firms? Only if we believe the market has (1) made a mistake and (2) will figure it out soon. Amazon’s high P/E ratio might reflect an overly optimistic assessment of its future growth or conservative estimate of its risk exposure, in which case it might be time sell Amazon stock if we own it or short it if we don’t. But, notice that we are betting that we have better information than the market (i.e., thousands of professional investors) *and* the market will eventually figure out we’re correct. If the market doesn’t figure out its error, we might miss out on further upside if we sold, or be forced to cover a lot of margin calls if we sold-short.

A few things to keep in mind. First, negative P/E ratios are meaningless; investors don’t pay to lose money. If earnings are negative, the P/E ratio is undefined. Second, be careful of interpreting very large P/E ratios (e.g., 500). They may just reflect very low earnings and dividing by a small number can produce a very large number. For example, at the end of 2020, Tesla (TSLA), the electric car manufacturer, had a P/E ratio over 1,000 on \$0.21 of

earnings per share.

8.6.6 The PEG Ratio: P/E and Growth

When we consider the three determinants of P/E ratios - d , r , and g from equation 8.18 - the latter two are responsible for most of the variation across firms. Payout ratios vary less across firms and over time than expected returns and earnings growth rates. Therefore, firms tend to have different P/E ratios primarily because they have different expected returns or earnings growth rates.

These differences make comparing P/E ratios across firms more difficult. For example, is Amazon's P/E ratio higher than Microsoft's because Amazon has a lower expected return or greater earnings growth? It would be nice if we could hold fixed one of these determinants and then compare P/E ratios. One attempt to do this is the **PEG ratio**, computed as the P/E ratio divided by the estimated earnings growth rate.

$$\text{PEG Ratio} = \frac{\text{P/E Ratio}}{\text{Earnings Growth Rate (measured in \%)}} \quad (8.19)$$

There are several ways to measure the PEG ratio, which differ in how earnings growth is measured. We could use an average historical growth rate, last year's growth rate, next year's forecasted growth rate, or a longer-term annual forecasted growth rate, such as over the next five years.

Let's compute the 2022 PEG ratios for Microsoft and Amazon, whose earnings are forecasted to grow next year by 15.0% and 48.8%, respectively.

$$\begin{aligned} \text{MSFT PEG Ratio} &= \frac{32.3}{15} = 2.15 \\ \text{AMZN PEG Ratio} &= \frac{59.9}{48.8} = 1.23 \end{aligned}$$

Loosely speaking, the PEG ratio is telling us the price stock market participants are willing to pay for \$1 of earnings and 1% of earnings growth. Despite having a significantly higher P/E ratio, Amazon has a lower PEG ratio because its forecasted earnings growth is so large, 48.8% next year. While the P/E ratio suggests that the market is paying more for a dollar of earnings from Amazon than Microsoft, the PEG ratio suggests that this higher price is because of the greater earnings growth. Relative to Microsoft, the PEG ratio suggests Amazon is cheap.

A rule of thumb uses 1.0 as a benchmark for PEG ratios to determine over- and under-valuations. Companies with PEG ratios above one are deemed overvalued, and below one

are deemed undervalued. While convenient, this rule is ad hoc, and there is no compelling empirical evidence suggesting that the PEG ratio can be used to reliably forecast future stock returns.

8.6.7 Valuing Microsoft with Multistage Dividend Growth Models

The Gordon Growth Model struggles valuing firms that are in a period of high growth because of its assumption of constant growth. Let's relax this assumption by using what we learned earlier about the effect of investment horizon on value. Recall equation 8.7, which is repeated here but after replacing $Price_T$ with $Value_T$.

$$Value_0 = \underbrace{\frac{Div_1}{1+r} + \frac{Div_2}{(1+r)^2} + \dots + \frac{Div_T}{(1+r)^T}}_{\text{Short-run value}} + \underbrace{\frac{Value_T}{(1+r)^T}}_{\text{Long-run value}} \quad (8.20)$$

This expression says that the value of a stock today can be viewed as the sum of two components: (1) short-run value consisting of cash flows over a finite horizon, T , and (2) long-run value consisting of cash flows over the indefinite future after period T . This long-run value, $Value_T$ goes by several names including **terminal value** and **continuation value**.

By viewing stock value through the lens of equation 8.20, we can allow for an arbitrary number of dividend growth rates over the short-term; hence, the term **multistage growth model**. Let's return to Microsoft to illustrate.

Based on their return on equity (34.65%) and retention rate (71.88%), Microsoft's earnings are forecast to grow $0.3465 \times 0.7188 = 24.90\%$ over the next year. While they can't grow at that rate forever, they can grow at that rate for say the next 10 years. After 10 years, we'll assume that the Microsoft's growth has settled down to its long-run equilibrium level, which is dictated by the growth of the overall economy and was discussed earlier in the context of table 8. Let's take a conservative approach and estimate Microsoft's sustainable annual growth after 10 years at 4%. The two different growth rates imply a **two-stage dividend growth model**.

Figure 8.11 presents the timeline of cash flows assuming a forecasted dividend of \$2.81 per share and a discount rate of 8.47%. Dividends grow at 24.9% for the next 10 years. From year 11 onward, dividends grow at a constant 4%. The forecasted stock price 10 years from today is the terminal value, which is a growing perpetuity.

$$Price_{10} = \frac{D_{11}}{r-g} = \frac{D_{10} \times (1+g)}{r-g} = \frac{20.786(1+0.04)}{0.0847-0.04} = \$483.61$$

Year	0	1	2	3	9	10
Dividends		2.81	$2.81(1.249)$	$2.81(1.249)^2$	$2.81(1.249)^8$	$2.81(1.249)^9$
Terminal value						483.61
Cash Flows		2.81	3.509	4.383	16.649	504.60

Figure 8.11: Timeline for Two-Stage Dividend Discount Model of Microsoft




Figure 8.12: Dividend Growth Paths of Two-Stage and Three-Stage Dividend Growth Models of Microsoft

Discounting all of the cash flows and summing produces a current stock price of \$267.56 per share - still low relative to the market price of \$311.21 but much closer than values from the Gordon Growth Model.

One troubling aspect with our two-stage growth model is the abrupt shift from high to low growth. Ten years from today, dividend growth declines from over 24% to 4% within a year. A more realistic assumption is that dividend growth gradually declines to its long-run equilibrium level of 4%. For example, assume Microsoft's dividend grows at 24.9% for the next five years, then that rate of growth declines by 2% per year for the following ten years before reaching 4% in years 16 and beyond. Figure 8.12 illustrates the two dividend growth paths.

The timeline for our three-stage model is presented in figure 8.13. We've broken the time line into segments corresponding to the different growth regimes.

Valuing the stock requires discounting all of these cash flows by the equity cost of capital.




Figure 8.13: Timeline for Three-Stage Dividend Discount Model of Microsoft

The value of the stock 15 years from today is the value of the growing perpetuity.

$$Value_{15} = \frac{Div_{16}}{r - g} = \frac{Div_{15}(1 + g)}{r - g} = \frac{24.81(1 + 0.04)}{0.0847 - 0.04} = \$577.34$$

Discounted back today, we get $577.34 \div (1 + 0.0847)^{15} = \170.53 . The present values of the other cash flows are calculated in the usual manner. Summing them produces a stock price estimate of \$257.01.

8.6.8 How Many Stages of Growth?

This discussion begs a couple of practical questions. How many growth stages should a company have? And, when should sustainable growth be assumed to kick in? In other words, at what point do we throw our hands up and just assume perpetual growth? Answers to these questions depend on several factors including where the company is in its life cycle, technological innovation in its industry, and its specific growth prospects. A young firm in a high growth sector may take 20 years, if not more, before settling down into sustainable growth. Consider Microsoft. After over 50 years of business, it is still a “growth” company, though it struggled with slow growth for a number of years after 2000.

As a practical matter, many financial analysts will forecast annual growth rates up to some time T , after which they will rely on a sustainable growth rate. Specifically, the pattern of earnings growth from today to time T can be anything, varying every year if necessary. As the terminal period is approached, earnings growth will begin to settle down

to its sustainable level. Once this transition is made, sustainable growth of nominal cash flows should not exceed 7%; with real cash flows, sustainable growth should not exceed 5% (and these estimates are on the *high* side of defensible). Of course, sustainable growth can be *less* than these estimates, including even negative. Many, businesses will eventually contract because of new technologies, competition, or poor management (e.g., print media in the early 2000s).

In sum, multi-stage dividend discount models provide additional flexibility by allowing for different growth regimes but are still limited by their assumption that shareholders only receive cash flows in the form of dividends.

8.7 Total Payout Model

In the dividend discount model, we emphasized the importance of accurately estimating dividend growth and, implicitly, the equity cost of capital because value is very sensitive to their difference, $r - g$. Additionally, next period's dividend is relatively easier to forecast.

However, consider table 12, which presents Microsoft's shareholder distributions and issuances, as well as their earnings and estimated payout ratios. These data are from Microsoft's income and cash flow statements, which have a fiscal year end of June 30. For every dollar of dividends Microsoft has paid since 2019, it has averaged \$1.58 of share buybacks. Its **total payout ratio**, which is the fraction of earnings distributed in dividends or share repurchases, is more than double the typical payout ratio. Ignoring repurchases significantly understates the cash flows going to Microsoft's shareholders.

	2023 (TTM)	2022	2021	2020	2019
Earnings, <i>Earn</i>	67.45	72.74	61.27	44.28	39.24
Dividends, <i>Div</i>	18.96	18.14	16.52	15.14	13.81
Repurchases, <i>Repo</i>	28.61	32.70	27.39	22.97	19.54
Issuances, <i>Iss</i>	1.76	1.84	1.69	1.34	1.14
Total payout, $Div + Repo$	47.58	50.83	43.91	38.11	33.35
Total net payout, $Div + Repo - Iss$	45.82	48.99	42.21	36.76	32.21
Payout ratio, $Div/Earn$ (%)	28.12	24.93	26.96	34.18	35.20
Total payout ratio, $(Div + Repo)/Earn$ (%)	70.53	69.88	71.66	86.05	85.00
Total net payout ratio, $(Div + Repo - Iss)/Earn$ (%)	67.93	67.35	68.90	83.02	82.09

Table 12: Microsoft Shareholder Cash Flows (\$million)

Figure 8.14 shows that Microsoft is not alone in its use of stock buybacks to return capital to shareholders. The bars show that the total dollars spent on repurchases grew more rapidly




Figure 8.14: Aggregate Dividends and Buybacks Paid by U.S. Firms and the Percentage of Firms Paying a Dividend or Engaging in Share Buybacks. Source: Standard and Poor's Compustat Data

than, and eventually surpassed, the money spent on dividends between 2018 and 2022. The figure also shows that the percentage of firms paying dividends declined from over 50% to less than 30%, while the fraction of companies repurchasing shares increased from 23% to 36%. The growth of repurchases over this period is due to a combination of factors including legal protection for companies repurchasing their shares, differential tax treatment of capital gains and dividend income, and the flexibility repurchases afford companies.¹⁴

The prevalence and importance of share repurchases motivates our next stock valuation model - the **Total Payout Model** which can be expressed as follows.

$$Value_0 = \frac{d_1 \times Earn_1}{1+r} + \frac{d_2 \times Earn_2}{(1+r)^2} + \frac{d_3 \times Earn_3}{(1+r)^3} + \dots \quad (8.21)$$

This equation is just our Fundamental Value Relation in which the cash flows in period t are equal to the **total payout ratio**, d_t , which includes both dividends and repurchases, times

¹⁴SEC rule 10b-18 provides a safe harbor for firms repurchasing their shares as long as they abide to certain guidelines, such as purchasing shares through a single broker and purchasing less than 25% of the average daily volume. Repurchases are taxed at a lower rate capital gains rate - assuming the share is held for more than one year - than dividend income. The market expects gradual consistent increases in dividends, whereas repurchase programs come with no such expectation.

earnings or net income, $Earn_t$.¹⁵ Unlike the dividend discount model, we can't use earnings per share to estimate the stock price. Instead, we estimate the total value of equity (i.e., market capitalization) using total earnings from the income statement, and then divide by the current number of shares outstanding to get the stock price.

To implement the total payout model, we have to make some assumptions similar to the dividend discount model. For example, we could assume that the total payout ratio is constant over time so $d_t = d$ for all t , and that earnings grow at a constant rate, g . In this case, the total payout model is a growing perpetuity and looks very similar to the Gordon Growth model with the difference being the use of the *total* payout ratio (dividends plus repurchases divided by earnings) instead of the payout ratio (dividends divided by earnings).

$$\begin{aligned} Value_0 &= \frac{d \times Earn_1}{(1+r)} + \frac{d \times Earn_1(1+g)}{(1+r)^2} + \frac{d \times Earn_1(1+g)^2}{(1+r)^3} + \dots \\ &= \frac{d \times Earn_1}{r-g} \end{aligned} \quad (8.22)$$

Alternatively, we can construct a multistage total payout model that allows for different earnings growth over the short-run, and a constant sustainable growth beyond some point in time, T .

$$Value_0 = \underbrace{\frac{d_1 \times Earn_1}{(1+r)} + \frac{d_2 \times Earn_2}{(1+r)^2} + \dots + \frac{d_T \times Earn_T}{(1+r)^T}}_{\text{Short-run value}} + \underbrace{\frac{Value_T}{(1+r)^T}}_{\text{Long-run value}} \quad (8.23)$$

8.7.1 Valuing Microsoft (Again)

Let's take another crack at valuing Microsoft, this time using a multistage version of the Total Payout Model. We'll assume earnings grow at the analyst forecasted rate of 12.79% for the next 10 years, after which they'll grow at their sustainable rate of 4%. We'll also assume that Microsoft distributes 76.63% (their recent average) of their earnings to shareholders via dividends and repurchases, and that their equity cost of capital is 8.47%. The cash flows are displayed in figure 8.15.

Discounting the cash flows by an 8.47% equity cost of capital produces an estimated market capitalization of \$2,421.59 billion. Microsoft had 7.44 billion shares outstanding as of July 2023, implying an estimate stock price of \$325.44 per share, slightly above the current

¹⁵A modification to the total payout model recognizes that firms receive money from shareholders via issuances. **Total net payout** is therefore dividends plus repurchases minus issuances, in which case d_t is the fraction of earnings distributed to shareholders net of any issuances.

Year	0	1	2	3	9	10
Earnings		76.08	85.81	96.78	199.26	224.74
Payout		58.29	65.75	74.16	152.68	172.21
Terminal value						$\frac{172.21(1+0.04)}{0.0847-0.04}$
Cash Flows		58.29	65.75	74.16	152.68	4,178.69

Figure 8.15: Timeline for Two-Stage Total Payout Model of Microsoft

market price of \$311.21. To avoid the sharp decline in earnings growth 10 years from today, we can employ a three stage model with a gradual transition to sustainable growth, but let's leave that as an exercise.

8.8 Flow to Equity Method

The last equity valuation model we'll discuss bears a lot of similarities to the discounted cash flow analysis we performed for capital budgeting in chapter 5. The **flow to equity** method differs from the dividend discount and total payout models only in how it measures cash flows. Rather than using dividends or total payouts, the flow to equity method uses **levered free cash flows** or **free cash flow to equity (FCFE)**.

$$Value_0 = \frac{FCFE_1}{1+r} + \frac{FCFE_2}{(1+r)^2} + \frac{FCFE_3}{(1+r)^3} + \dots \quad (8.24)$$

Levered free cash flow is defined as

$$\begin{aligned}
 FCFE = & \overbrace{(Sales - Expenses - D\&A) \times (1 - \tau) + D\&A - NLT I - NWCI}^{\text{Unlevered free cash flow (FCF)}} \\
 & - \underbrace{Interest + NDI}_{\text{Cash flow to creditors}} \\
 & + \underbrace{\tau \times Interest}_{\text{Interest tax shield}}, \quad (8.25)
 \end{aligned}$$

where *Interest* is the interest expense on any borrowings, τ is the effective tax rate, and *NDI* is net debt issuance (debt issuances minus debt repurchases).¹⁶

¹⁶Since the passage of the Tax Cuts and Jobs Act in 2018, Equation 8.25 is only an approximation. The Act limits the amount of interest expense a company can deduct to 30% of their taxable income. Any additional interest expense can be carried forward and applied against future tax liabilities.

Specifically, we want to choose a short-term forecast horizon, T , so that the growth of our cash flows in the last year, from $T - 1$ to T , is close to g . “Close” is subjective, but consider the examples in Table 13.

	Short-term forecast horizon				
	1	2	3	4	5
Cash flows 1	80.0	98.4	114.1	124.4	126.9
Year-on-year (YoY) growth		23.0%	16.0%	9.0%	2.0%
Cash flows 2	500.0	525.0	551.3	578.8	607.8
Year-on-year (YoY) growth		5.0%	5.0%	5.0%	5.0%
Cash flows 3	250.0	250.0	237.5	213.8	181.7
Year-on-year (YoY) growth		0.0%	(5.0%)	(10.0%)	(15.0%)
Cash flows 4	10.0	20.0	50.0	150.0	525.0
Year-on-year (YoY) growth		100.0%	150.0%	200.0%	250.0%

Table 13: Example Cash Flow Growth Profiles

Table 13 presents three sets of cash flow forecasts for a short-term forecast horizon of 5 years. Cash flows 1 and 2 have terminal year growth rates of 2% and 5%, respectively. Both of these growth rates offer a natural transition to a terminal rate around 3%. Cash flows 3 are contracting over time suggesting that this business is in decline. We could continue to forecast cash flows to some terminal point in time, like bankruptcy or liquidation, or we could assume a negative growth rate for our terminal value. Finally, cash flows 4 exhibit increasing growth consistent with a small or early-stage firm. Assuming that growth suddenly slows from 250% in year 5 to 3% in year 6 makes little sense in this case. A longer short-term forecast horizon, 15 to 20 years, in this example makes more sense.

Sometimes people will argue: “We can’t be confident in our forecasts too far into the future so it’s better to choose a small T , like 3 or 5 years.” This is a nonsensical argument. We’re making a forecast for *every* year regardless of what we assume for T ! What’s relevant is that cash flow growth stabilize by the time we get to T .

A word of warning: cash flows can be highly volatile, especially for high growth firms. Young or fast growing firms can have volatile cash flows because they are investing heavily. As such, recent cash flow growth can be quite low, even negative. In these and other cases in which the firm is undergoing significant investment, earnings or even sales growth may be more informative about the current and future growth potential of the company.

8.9 Key ideas

Stocks, like bonds, are an important financial instrument. They are a central source of funding for companies and an important investment for savers.

- Shareholders own the company. The size of each shareholder's ownership stake is proportional to the number of shares they own with single class structures and one vote per share.
- The three most common stock valuation models are just applications of our Fundamental Value Relation with different ways of estimating the cash flows that shareholders receive.
 1. Dividend discount model (DDM),
 2. Total payout model (TPM), and
 3. Flow to equity model (FTE).
- Within each model, we could vary our assumptions about the growth of cash flows. Constant growth implies each model is a growing perpetuity, which while simple and informative, has limited application to companies not operating at their sustainable growth rate. Multistage growth models relax this assumption by allowing for more than one era of growth rate.
- The sustainable growth rate that persists indefinitely into the future is capped by the growth rate of the overall economy - 7% in nominal terms, 4% in real terms. When firm's reach their sustainable growth rate determines the "cap T ," or the time in our model at which high growth subsides and sustainable growth takes over.
- The expected return on a stock is equal to the dividend yield plus the price appreciation/depreciation (a.k.a., capital gain/loss). Price appreciation/depreciation is in turn equal to dividend or earnings growth according to the dividend discount and total payout models, respectively. For stock prices to increase, dividends or earnings must grow.
- P/E ratios are a positive function of how much the firm distributes to shareholders (payout ratio, d) and earnings growth (g), and a negative function of expected returns (r). In other words, *all else equal*
 1. Paying more to your shareholders will increase the P/E ratio.

2. Increasing the growth of earnings will increase the P/E ratio.
3. Increasing the expected return on the stock will decrease the P/E ratio.

The “all else equal” caveat is important because changing any one of the three P/E ratio drivers will typically change the other two.

- To make the P/E ratio more easily comparable across firms with different growth rates, the PEG ratio will divide the P/E ratio by the estimate growth rate. However, this does not solve the problem of different capital intensities, financial policies, and tax strategies confounding any cross-company comparisons.
- We can use any of our models to estimate expected stock returns, or we can use historical data and statistical averages. But, we must not lose sight of the volatility of stock returns - a measure of risk - which can also be estimated with historical averages.

8.10 Technical Appendix

8.10.1 ROE and Retention Rate

Equation 8.12 estimates earnings growth as the product of the return on equity and the retention rate - $ROE \times (1 - d)$. Assuming all earnings growth comes from new investment, as opposed to existing assets, and firms rely on earnings, as opposed to external capital markets, this relation can be derived as follows.

$$\begin{aligned} Earn_{t+1} - Earn_t &= ROE_t \times Investment_t \\ &= ROE_t \times Earn_t \times (1 - d_t) \end{aligned}$$

where $Earn_s$ are the earnings for period s . Dividing both sides by $Earn_t$ produces the result.

$$\begin{aligned} \frac{Earn_{t+1} - Earn_t}{Earn_t} &= \frac{ROE \times E_t \times (1 - d)}{E_t} \\ g &= ROE \times (1 - d) \end{aligned}$$

8.10.2 PVGO and Profitable Investment

Equation 8.17 decomposed today's stock price into two components

$$Price_0 = \underbrace{\frac{Earn_1}{r}}_{\text{No-growth value}} + \left(\underbrace{\frac{d \times Earn_1}{r-g} - \frac{Earn_1}{r}}_{\text{Value from growth}} \right)$$

The second term, the value from growth, is positive if the return on investment (ROE) is greater than the equity cost of capital, r .

$$\begin{aligned} \frac{d \times Earn_1}{r-g} - \frac{Earn_1}{r} > 0 &\iff \frac{d \times Earn_1}{r-g} > \frac{Earn_1}{r} \\ &\iff d > \frac{r-g}{r} \\ &\iff \frac{g}{r} > 1-d \\ &\iff \frac{(1-d)ROE}{r} > 1-d \\ &\iff ROE > r \end{aligned}$$

This result emphasizes what was mentioned earlier. Growth for growth's sake is not useful and can be value destructive if the investment return generating the growth is less than cost of capital - what was previously discussed in chapter 5.

8.10.3 Taxes

Like bonds, earnings from stock investments are taxed in the U.S. Specifically, capital gains and dividend income are both taxed, but typically at different rates. Dividends are taxed as ordinary income and as such depend the tax bracket into which we fall - higher income, higher tax rate. Capital gains are taxed similarly, which higher income investors taxed at a higher rate. Table 14 presents the capital gains tax rates as of 2021, which are significantly higher than the corresponding income tax rates.

There is one important wrinkle to gains taxes. To take advantage of the lower tax rate, the capital gains must be deemed **long-term capital gains**, which in practice means the gain must be experienced over at least one year. In other words, we have to buy and hold the stock (or bond) for at least one year to receive the beneficial tax treatment. If the stock is sold within one year, the any capital gains are treated as ordinary income and tax accordingly.

Tax-filing status	0% tax rate	15% tax rate	20% tax rate
Single	\$0 to \$41,675	\$41,676 to \$459,750	\$459,751 or more
Married, filing jointly	\$0 to \$83,350	\$83,351 to \$517,200	\$517,201 or more
Married, filing separately	\$0 to \$41,675	\$41,676 to \$258,600	\$258,601 or more
Head of household	\$0 to \$55,800	\$55,801 to \$488,500	\$488,501 or more.

Table 14: 2021 Capital Gains Tax Rates by Filing Status and Income Bracket

8.11 Problems

8.1 (*Ownership structure*) Who owns a company: Shareholders or Creditors?

8.2 (*Ownership structure*) Chewey's is a retail pet store with 5 million shares outstanding.

Using this information, answer the following questions.

- If Bandit owns 2.501 million shares and Luna owns the rest, who has control of the company assuming each share has one vote? That is, who is the ultimate decision maker?
- If Bandit and Luna both owned 2.5 million shares, who would be in control of the company assuming each share has one vote?
- If Bandit, Luna, and French Fry each owned one third of the outstanding shares, who would be in control of the company assuming each share has one vote?
- If Bandit, Luna, and French Fry each own 1 million shares, and Gumby owned the remaining 2 million, who would be in control of the company assuming each share has one vote?

8.3 (*Multi-class ownership structures*) 1-800-FLOWERS.COM is an e-commerce company that began as a teleflorist. The company has a dual class share structure. Class A shares are traded on the NASDAQ and have 1 vote per share, class B shares do not trade on public markets and have 10 votes per share. As of September 2022, there were 37,287,993 class A shares and 27,249,614 class B shares outstanding, which can be converted into class A shares on a one-for-one basis.

Using this information, answer the following questions.

- How are voting rights split between class A and class B shares? What fraction of the votes does each class hold?

- b. Why would class B shareholders convert their stock to class A?
- c. How many shares can class B shareholders convert to class A to and still maintain control of the company?

8.4 (*Share issuance and ownership*) Xarc Technologies is a founder-run company specializing in logistics software solutions. The sole owner, Paulina Hostin, currently owns 10 million shares estimated to be worth 50 million dollars. She is considering raising new capital (i.e., money) by issuing shares to the public.

Using this information answer the following questions.

- a. What is the current price per share?
 - b. Assuming Paulina issues shares at the current market price, each with one vote like her shares, how many shares can she issue and still retain control of the company? How much money can she raise?
 - c. If she needs to raise \$100 million for investment, how many shares must she issue at the current price? What will her ownership percentage be after the issuance? Will she still be in control of the company?
 - d. How can Paulina restructure the voting rights - number of votes per share - of her shares so she can raise \$100 million in equity capital and still retain control?
- 8.5 (*Dividends*) True or False. Like interest for creditors, firms must pay their shareholders dividends or risk legal action.
- 8.6 (*Priority structure*) In what order are bondholders, common stockholders, and preferred stockholders paid? That is, what is the priority structure of these claims?
- 8.7 (*Preferred dividends*) Triton International Limited is a leasing company specializing in intermodal freight equipment leasing and maritime container management services. They have outstanding cumulative preferred shares with a face value \$25 and coupon rate of 6.875% (ticker TRTN-D).

Using this information, answer the following questions.

- a. What is the annual dividend on the preferred shares?
- b. What is the quarterly dividend on the preferred shares?
- c. If Triton misses a dividend payment to the holders of TRTN-D shares, does it have to make it up before paying dividends to common shareholders?

d. The share price of TRTN-D was]\$23.83. What is the corresponding dividend yield?

8.8 (*Stock quotes*) Alex Barnow was interested in purchasing some shares in Shake Shack, an American fast casual restaurant chain specializing in hamburgers. His trading platform showed the following information as of 7:00 PM April 21, 2023.

Last sale x size	Change
55.69 x 39,881	+ 0.09 (+0.16%)
Bid x size	Ask x size
42.00 x 1	60.00 x 4

Using this information, answer the following questions.

- To what do the numbers under “Last sale x size” refer?
- To what do the numbers under “Bid x size” refer?
- To what do the numbers under “Ask x size” refer?
- To what do the numbers under “Change” refer?

8.9 (*Stock orders*) Pfizer Inc. (PFE) is a large pharmaceutical company trading on the New York Stock Exchange. As of April 21, 2023, the bid and ask quotes, and corresponding trade sizes, reveal the following information.

Bid x size	Ask x size
40.11 x 4,000	40.35 x 4,000

Using this information, answer the following questions.

- A market buy order for 100 PFE shares is likely to execute at a price near what number? What risk do we face when submitting this order?
- If we want to sell 50 shares of PFE at a price no lower than \$40 per share, what type of order should we submit? (Be specific.) What risk do we face when submitting this order?
- If we submit a limit buy order at \$27 per share, what risk are we taking?
- If we want to purchase 20,000 shares of PFE and ensure that all shares are purchased at the same price, what type of order should we submit?

- 8.10 (*Stock orders*) Martha Stewart claims to have submitted to her broker, Peter Bacanovic, a stop-loss order at \$60 per share for all of her Imclone Systems Inc. (IMCL) stock holdings. What does this order mean?
- 8.11 (*Stock orders*) Price quotes for Johnson and Johnson (JNJ) as of July 20, 2023 are show in the following table.

Bid x size	Ask x size
168.21 x 1,000	168.22 x 2,000

What type of order should we submit if we want to...

- buy 500 shares of JNJ as quickly as possible.
 - sell 100 shares of JNJ stock at no less than \$167 per share.
 - buy 10,000 shares of JNJ stock at no more than \$169 per share and ensure that all 10,000 shares are purchased at once.
 - sell 200 shares if the price of JNJ falls below \$160 per share.
 - buy 50 shares if the price reaches \$170 per share.
- 8.12 (*Market capitalization, ownership, realized returns*) In March of 2022, Sam Kingston purchased 500 shares of AAR Corp. (AAR), an aerospace and defense contractor, at a price of \$48.85 per share. At the time, AAR had 28.985 million shares outstanding. Using this information, answer the following questions.
- What was the market capitalization of AAR at the time of Sam's purchase?
 - What fraction of the company did Sam purchase?
 - If we are only told that the price per share of AAR was \$50.64 in March 2023, do we have enough information to compute Sam's total return for the year? Price return?
 - If we are only told that Sam's total annual return on her investment in AAR was 4.86%, can we determine the price at which she sold her shares?
- 8.13 (*Realized returns*) Toby Howard purchased 10,000 shares of Ring Energy (REI) for \$10.25 per share in August of 2015. Toby sold his shares in August 2022 for \$17.76 per share. What is Toby's total return between 2015 and 2022 assuming Ring paid no dividends and did not split its shares? What is his annualized return?

8.14 (*Realized return, dividend yield*) The table below presents market data for Goldman Sachs (GS). The Price column presents the ex-dividend closing price, i.e., the price of the day's last trade after the dividend payment (or as of the ex-dividend date).

Date	Dividend (\$/share)	Price (\$/share)
8/30/2021	0.00	413.60
8/31/2021	2.00	413.51
9/1/2021	0.00	413.66

Using this data, answer the following questions.

- What are the one-day price returns?
- What are the one-day dividend yields?
- What are the one-day total returns?
- If an investor purchased one share of Goldman Sachs stock at the closing price on 8/31/2021 and sold his stock at the closing price on 9/1/2021, what would his total return be ignoring any transaction costs? How would this return change if he purchased 100 shares of Goldman stock instead of one?
- What is the two-day total return to an investor that purchases Goldman Sachs stock on 8/30/2021 at the closing price and sells on 9/1/2021 at the closing price assuming he receives the dividend?

8.15 (*Realized returns, volatility*) The table below presents market data for GameStop (GME). The Price column presents the closing price, i.e., the price of the day's last trade. The Volume column presents the total number of shares traded during the week. GameStop paid no dividend during this time.

Date	Price (\$/share)	Volume (millions of shares)
1/8/21	17.69	33.65
1/15/21	35.50	307.07
1/22/21	65.01	409.30
1/29/21	325.00	559.24
2/5/21	63.77	302.04
2/12/21	52.40	116.62
2/19/21	40.59	70.83

Using this data, answer the following questions.

- a. What were the realized weekly returns from January 15 to February 19?
- b. Estimate GameStop's expected weekly return by taking the arithmetic average of the returns computed for the previous question. What is the corresponding annual expected return?
- c. Estimate GameStop's weekly volatility with the standard deviation of the the returns computed for the first question. What is the corresponding annual volatility?
- d. What is the growth rate of trading volume from January 8th to January 29?
- e. If you had sold short 100 shares of GameStop stock at the closing price on January 22nd, how much money would you have received ignoring any transaction costs? If you were forced to buy and return the stock on January 29th at the closing price because of a short squeeze, how much money would you have to spend? What would your net gain/loss be from this transaction?

8.16 (*Realized returns, expected returns, volatility*) The table below presents market data for Tesla (TSLA). The Price column presents the closing price, i.e., the price of the day's last trade. Tesla paid no dividends during 2021.

Date	Price
12/30/2020	705.67
02/29/2021	793.53
02/26/2021	675.50
03/31/2021	667.93
04/30/2021	709.44
05/28/2021	625.22
06/30/2021	679.70
07/30/2021	687.20
08/31/2021	735.72
09/30/2021	775.48
10/29/2021	1,114.00
11/30/2021	1,144.76
12/31/2021	1,056.78

Using this data, answer the following questions.

- a. What were the realized monthly returns for 2021.

- b. Using the returns computed in the previous problem, estimate the monthly expected return using the arithmetic average. What is the corresponding annual expected return?
- c. Using the returns computed in the first problem, estimate the monthly return volatility using the standard deviation. What is the corresponding annual volatility?
- d. What was the realized annual return for 2021? How does it compare to the expected return estimated in question b.?

8.17 (*Expected returns, dividend growth*) We have decided to purchase 250 shares of Visa stock on May 30, 2008 for \$86.36 per share. Visa currently has announced a quarterly dividend of \$0.105 per share payable on August 29, 2008. Visa maintains a constant dividend for four quarters and then increases it by 20%. So, Visa will pay a \$0.105 dividend per share at the end of August 2008, November 2008, February 2009, and May 2009 after which the dividend will then grow by 20% but remain constant for four more quarters. We anticipate selling our shares on February 28, 2012 for \$116.37 per share. (Treat all quarters as being equal length.)

Using this information, answer the following questions.

- a. What is our investment expense on May 30, 2008?
- b. What are the dividends per share over our investment horizon? Construct a timeline.
- c. What is the price return of our investment over the entire investment horizon? What is the annualized return?
- d. What are the quarterly and annual internal rates of return of our investment accounting for the dividends? Do these estimates accurately reflect our realized returns? Explain why or why not.
- e. Using your answers to the previous questions, estimate the annual dividend of our investment.

8.18 (*Realized returns, dividend reinvestment*) Monthly closing prices and dividends per share for Dominion Energy Inc. (D) are presented in the following table.

Date	Price	Dividend
5/31/22	84.22	
6/30/22	79.81	0.6675
7/29/22	81.98	

Using this data, answer the following questions.

- a. What are the monthly returns for June and July?
- b. What is the two-month return from May 31 to July 29? What are the corresponding monthly compounded and annual compounded returns?
- c. Show that the two-month return corresponds to a strategy of buying the stock at the end of May for \$84.22, reinvesting the dividend in the stock for one month, and then selling the stock for \$81.98 at the end of July.
- d. What is the monthly internal rate of return on an investment in Dominion that purchases stock on May 31 at the closing price and sells on July 29th for the closing price? How does this compare to the monthly compounded return from question b.? Explain the difference.
- e. Show that the IRR estimates the monthly return to a strategy in which we buy the stock at \$84.22 at the end of May, reinvest the dividend in another investment earning the IRR for one month, and then sell the stock for \$81.98 at the end of July.
- f. If we purchase 1,000 Dominion shares at the end of May, how many shares will we own at the end of July assuming we reinvest all dividend into the stock? (Fractional share ownership is allowed.)

8.19 (*Mispricing and returns, dividend discount model*) Realty Income Corporation (O) is a real estate investment trust (REIT) that owns a portfolio of commercial properties. As of July 23, 2023, Realty Income's stock price was \$63.03 and it is forecast to pay a \$3.07 dividend per share next year. Their equity cost of capital is 8.00%.

Using this information, answer the following questions.

- a. What is the forward dividend yield?
- b. What would you predict for next year's price?
- c. What the predicted price return?
- d. What is the market implied (i.e., using the market price) dividend growth rate according to the Gordon growth model? How does it compare to your answer to the previous question?
- e. What does the sum of the forward dividend yield and dividend growth rate equal?

8.20 (*Dividend distribution timing*) On Thursday September 15, 2022 Abbot Laboratories announced a \$0.47 dividend per share to be paid on Tuesday November 15, 2022. The record date for the dividend was Friday October 14, 2022.

Using this information, answer the following questions.

- What is the declaration date?
- What is the payment date?
- When is the ex-dividend date?
- By what date must an investors own shares in Abbott in order to receive the dividend?

8.21 (*Realized returns, stock splits*) The table below presents daily stock market data for Amazon (AMZN).

Date	Dividend (\$/share)	Price (\$/share)	Shares Outstanding (000s)
6/2/22	0	2,510.22	508,720
6/3/22	0	2,447.00	508,720
6/6/22	0	124.79	10,174,400
6/7/22	0	123.00	10,174,400
6/8/22	0	121.18	10,174,400

Using this data, answer the following questions.

- What was the split ratio for the stock split Amazon executed on June 6th?
- What was the June 6th return on Amazon stock?
- Did Amazon shareholders make or lose money on June 6th?

8.22 (*Expected return decomposition*) The table below contains information on expected equity returns, risk-free rates measured by the yield on 1-year T-bills, and risk premia. Fill in the cells containing letters with the appropriate estimate.

Company	Risk-free rate (%)	Risk premia (%)	Expected return (%)
Blackstone (BX)	4.95	5.1	a.
KKR (KKR)	b.	4.74	9.32
Coterra Energy (CTRA)	4.74	c.	6.18
Altria Group (MO)	5.1	2.85	d.

8.23 (*Comparing returns, annualized returns, volatility*) Return summary statistics for four different stocks are presented in the following table. The sample period is the time frame over which the statistics are computed using monthly return data. The cumulative return is cumulative monthly compounded return from the start to the end of the sample period.

Stock	Sample Period	Average Monthly Return (%)	Monthly Volatility (%)	Cumulative Return (%)
Southern Co. (SO)	1/2005 - 12/2022	0.83	4.34	388.38
3M Co. (MMM)	6/2012 - 12/2022	0.68	5.68	92.62
Lilly Eli & Co. (LLT)	10/2018 - 12/2022	2.88	7.90	267.57
Advanced Micro Devices Inc. (AMD)	6/2014 - 12/2022	4.10	16.84	1,519.25

Using this information, answer the following questions.

- What is the estimated annual expected return for each stock?
- What is the ranking of stocks in terms of return performance?
- What is the ranking of stocks in terms of volatility?
- How do your answers to the previous two question relate?
- For each stock, compute the ratio of (i) the estimated annual expected return to (ii) the estimated annual volatility. Rank stocks by this ratio. How does this ranking compare to the rankings by expected return and volatility? How can we interpret this ratio?

8.24 (*Investment horizon, activist investor, dividend discount model*) Bichtel Manufacturing is planning to pay a dividend of \$1.84 per share one year from today, which it expects will grow by 4.8% per annum indefinitely under their current operating strategy.

Grim Reaper Investors is an activist investor group that has targeted Bichtel for what it perceives are poor governance practices. Grim Reaper is proposing to increase Bichtel's dividend to \$10 per share over the next two years in an effort to reduce their cash holdings and cut down on management's wasteful spending. Three years from today, Bichtel's dividend will return to its previously forecast level as of three years from today. However, sustainable dividend growth will be 0.5% less than what analysts were expecting. Grim Reaper anticipates no change to the equity cost of capital, which is 12.60%.

Using this information, answer the following questions.

- a. What are the dividend per share forecasts for the next four years under Bichtel's current operating strategy?
- b. What would you estimate as Bichtel's price per share if they continue with their current operating strategy?
- c. What are the dividend per share forecasts for the next four years under Grim Reaper's proposed strategy?
- d. What would you estimate as Bichtel's price per share if they implement Grim Reaper's proposal?
- e. Based on your answers to the previous questions, does Grim Reaper's proposal create or destroy value for Bichtel shareholders and by how much?
- f. If Grim Reaper buys Bichtel's stock at the price reflecting Bichtel's current operating strategy, implements its strategy, and then exits two years from today after receiving dividend two years from today, what will be its exit price and annual return on investment? How does it compare to Bichtel's expected return to its current operating strategy?
- g. Assume for this question that Grim Reaper's proposal increases the estimated price per share of Bichtel to \$35.97. What sustainable growth rate would Bichtel management have to reach to match the increase in value and maintain their current projected dividend of \$1.84 per share. Would this be feasible?

8.25 (*Gordon growth model, stock valuation*) ConocoX is a supplier of light rail equipment. It forecasts a dividend per share of \$0.86, which is expected to grow at 12% per year thereafter. ConocoX's estimated expected equity return is 14.8%. Use the Gordon growth model to estimate its current share price

8.26 (*Gordon growth model, misvaluation*) As of April 2023, Yahoo! Finance reports the following information for the Ford Motor Company (F), an American automotive manufacturer. Ford's current share price is \$12.09 on 4.037 billion shares outstanding. The forecasted dividend for next year is \$0.60 per share, which is expected to *contract* at a rate of 9.86% per year.

Using this information, answer the following questions.

- a. What is Ford's current market capitalization?
- b. What is Ford's prospective dividend yield?
- c. What is Ford's expected return according to the Gordon growth model?

- d. Ford's CEO Jim Farley has argued that analysts have confused a short-term contraction in the company, as it pivots to electric vehicles, for long-term growth which he estimates to be 5%. If he is correct, what is the expected return on Ford's equity?

8.27 (*Present value of growth opportunities*) Historical market and accounting information for Alphabet Inc - formerly known as Google Inc. - is presented in the table below.

Date	Price per share (\$)	Earnings per share (\$)	Equity cost of capital (%)
12/2008	307.65	13.31	9.93
12/2020	1,752.64	58.61	5.59

Using this information, answer the following questions.

- What is the value of Alphabet's assets in place as of 2008 and 2020, using the perpetuity approach discussed in the chapter?
 - Using your answers to the previous question, what is the corresponding value of Alphabet's growth opportunities as of 2008 and 2020?
 - Using your answers to the previous questions, what fractions of the share price are attributable to assets in place and growth opportunities in 2008 and 2020? Have these fractions changed? If so, why?
- 8.28 (*Gordon growth model, misvaluation*) General Electric Company (GE) is an American industrial conglomerate. GE's share price in April 2023 is \$100.17 on 1.0903 billion shares outstanding. The forecasted dividend for next year is \$0.32 per share, which is expected to grow at a rate of 24.6% per year.

Using this information, answer the following questions.

- What is GE's current market capitalization?
- What is GE's prospective dividend yield?
- What is Ford's expected return according to the Gordon growth model?
- Vikki van Duyne, an equity analyst covering GE, argues that GE is currently overvalued because of excessive growth expectations. She believes a more accurate price per share is \$92. What is the growth rate implied by this share price, assuming the expected return is as you computed for the previous problem? Is her logic correct?

- e. Assuming GE's stock is overvalued as suggested by Vikki, what should the realized return be next year if the market recognizes its mistake? What is the corresponding price per share?
- f. Assuming GE's stock is overvalued as suggested by Vikki, what should the realized return be next year if the market doesn't realize it's mistake? And, what is the corresponding price per share?
- g. Is the dividend growth estimate sustainable, meaning reasonable for the indefinite future?

8.29 (*Dividend discount model, financial statements*) According to yahoo! finance, Miller-Knoll, Inc. (MLKN) researches, designs, manufactures, and distributes interior furnishings worldwide. Abbreviated versions of their financials as of July 21, 2023 and several projections are presented in the table below.

Income statement (\$mil)	TTM	Balance sheet (\$mil)	May 2023
Sales	4,231	Cash	223,500
COR	2,773	Receivables	363,500
Gross profit	1,458	Inventory	487,400
Operating expense	1,249	Other current assets	101,800
Operating income	209	Total current assets	1,176,200
Net nonop expense	108	Net PPE	952,200
Pretax income	101	Other assets	2,146,400
Taxes	32	Total assets	4,274,800
Net income	69		
		Payables	592,300
Market data		Current debt	110,500
Dividend per share (\$)	0.75	Total current liabilities	702,800
Price per share (\$)	17.63	Long term debt	1,758,800
Shares outstanding (million)	78.80	Other liabilities	273,000
		Stockholders' equity	1,540,200
Projections		Total liabilities & stockholders' equity	4,274,800
Dividend per share (\$)	0.75		
Earnings growth, next year (%)	21.30		

Using this information, answer the following questions concerning MillerKnoll.

- a. What is the current market cap?
- b. What is the current dividend yield?

- c. What is the current earnings per share (EPS)?
- d. What is the current price-to-earnings (P/E) ratio?
- e. What is the current payout ratio?
- f. What is the current plowback ratio?
- g. What is the current return on equity?
- h. Using answer to the previous two questions, what is the implied sustainable growth rate? Is it in fact sustainable and what does your estimate imply about the long-run potential of the company? Explain.
- i. What are the projected earnings and earnings per share, assuming the number of shares remain unchanged, for next year? What is the corresponding forward price-to-earnings ratio? How does it compare to the current price-to-earnings ratio?
- j. Assuming the payout ratio and shares outstanding remain unchanged, what are the projected dividend per share and forward dividend yield?
- k. Assuming the payout ratio remains constant, what is MillerKnoll's equity cost of capital implied by the Gordon Growth Model and your answers to the previous questions? Discuss the plausibility of your estimate?

8.30 (*Gordon growth model, stock valuation*) The Coca-Cola Company (KO), a beverage manufacturer and distributor has the following financial information as of April 2023 according to Yahoo! Finance.

- Forecasted annual dividend of \$1.84 per share
- Long-term (5-year) earnings growth rate of 5.81%
- 4.332 billion shares outstanding
- Current market price of \$63.96 per share

Using this information and the Gordon Growth model, answer the following questions.

- a. What is Coca-Cola's prospective (i.e., future) dividend yield?
- b. What is Coca-Cola's expected stock price appreciation?
- c. What is Coca-Cola's expected equity return, assuming it pays out a constant fraction of its earnings in dividends?
- d. (*Challenging*) What is the probability that Coca-Cola's stock return next year exactly equals what we calculated in the previous question?

- e. What is the current market capitalization of Coca-Cola?
- f. What is the expected market capitalization of Coca-Cola next year, assuming dividends are paid out at the end of the year?

8.31 (*Gordon growth model, expected return*) C3.ai (AI) is an enterprise artificial intelligence software company. As of April 2023, Yahoo! Finance reports that C3.ai has never paid a dividend and has no intention of paying one in the near future. It's current earnings per share (EPS) is -\$2.20 reflecting its recent losses. However, its market capitalization is \$2.251 billion.

Using this information, answer the following questions.

- a. Can we use the Gordon growth model to estimate the expected return for C3.ai?
- b. Is C3.ai a “counterexample” to our Fundamental Value Relation? In other words, is the fact that C3.ai has negative earnings and has never paid a dividend evidence that the Fundamental Value Relation is either wrong or only applicable to some firms?

8.32 (*Gordon growth model, dividend payout, return on equity*) Nike Inc. (NKE) is an American athletic apparel and footwear company. As of April 2023, Nike distributes 37.18% of their earnings to shareholders via dividends, a percentage investors anticipate remaining constant in the future. Their trailing twelve month earnings estimate is \$5.48 billion, and their most recent book equity figure is \$14.53 billion. Nike has 1.5367 billion shares outstanding.

Using this information, answer the following question.

- a. What Nike's return on equity?
- b. What Nike's implied dividend growth rate according to the Gordon growth model? Comment on the sustainability of this estimate.
- c. What is Nike's aggregate dividend, i.e., the sum of dividends paid to all shareholders?
- d. What is Nike's market capitalization?
- e. What is Nike's implied expected equity return?

8.33 (*Dividend discount model, bankruptcy*) Oh-Krike is an Australian media company in decline because of increased competition. Despite no intention of paying a dividend because of a lack of cash, the company's 2.4 million outstanding shares currently enjoy a price of \$1.62 per share. Oh-Krike's equity cost of capital is 12.5%.

Using this information, answer the following question.

- a. What is Oh-Krike's current market capitalization?
- b. What is the implied liquidation value of Oh-Krike one year from today if the company were to be dissolved and its assets sold to other companies?
- c. A financial consultant has suggested that Oh-Krike's liquidation value is effectively zero because most of its assets are people who will leave the company. If she is correct, what should Oh-Krike's share price be today?

8.34 (*Multistage growth dividend discount model*) Airbnb (ABNB) is on an online platform that matches consumers looking for a temporary residence with "hosts" willing to rent their residence. The company's current earnings per share based on the trailing twelve months of earnings is \$2.98. Analysts have forecast earnings growth of 14.60% over the next year, and 22% per annum over the following nine years. Airbnb's equity cost of capital is 16%.

The company has announced it will pay a dividend of \$1.50 per share starting one year from now.

Using this information, answer the following question.

- a. What are the projected earnings per share over the next 10 years?
- b. What is next year's payout ratio?
- c. Assuming the payout ratio remains constant for the indefinite future, what are the forecasted dividend per share estimates over the next 10 years? What are the corresponding dividend growth rates for years two through five?
- d. Assuming dividend growth slows to a sustainable 5% after the next 10 years, what would you estimate as the price per share of Airbnb 10 years from today? What about today?
- e. How does your estimate of Airbnb's price today compare to its actual price of \$148.77. Do you think Airbnb is over- or under-valued? Or, is something missing from your analysis?

Chapter 9

Investing: Portfolios

Fundamental value relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

A portfolio is a collection of assets - stocks, bonds, real estate, commodities, projects, etc. What happens to our risk and return when we own more than one asset? **Modern Portfolio Theory** or **MPT** tells us.

This chapter

- introduces the notion of a portfolio weight to represent the relative amount of an asset in a portfolio,
- shows how to compute realized and expected returns for portfolios, using the same concepts for individual assets,
- introduces some common portfolio weighting schemes and popular stock indices, such as the S&P 500.
- provides a conceptual framework - mean-variance analysis - for assessing different portfolios,
- introduces diversification and illustrates how the risk of a portfolio can be *less* than the sum of the risk of the assets in the portfolio,

- applies our fundamental value relation to answer several questions including:
 - What are the “best” portfolios in which to invest? How do they change over time?
 - How can we decide among the “best” portfolios in which to invest?
 - How can we measure risk-adjusted performance?
 - What role do (near) risk-free investments, like Treasury securities, play in our portfolios?

Warning. There are some...long formulas and somewhat heavy notation in this chapter. Don't fear. The underlying math is still arithmetic and is easily implemented on a calculator or in a spreadsheet. The key is not to lose sight of the intuition, which we'll emphasize throughout.

To keep the discussion manageable, we'll focus on portfolios of stocks but emphasize throughout that the discussion applies more broadly. In parentheses next to company names, we'll put the company's stock ticker symbol, e.g., Microsoft (MSFT).

Despite its somewhat technical nature, the message of this chapter is simple and important. Investing in different types of assets - large cap stocks, small cap stocks, stocks from different industries, bonds, real estate, etc. - can both increase returns and reduce risk. This is why some refer to diversification as the only free lunch in finance.

9.1 Portfolio Returns

9.1.1 Microsoft and Ferrari Portfolio

How does the risk and return of a portfolio of assets differ from that of an individual asset? Table 1 presents monthly stock price data for Microsoft (MSFT) and Ferrari (RACE). Specifically, the table presents for first and last few months of the period January 2017 to December 2021 the following data: share price at the end of the month, dividend per share, and monthly total realized return.

Recall from chapter 8 that the return on a stock over the period $t - 1$ to t , $r_{t-1,t}$, can be computed as follows:

$$r_{t-1,t} = \frac{Price_t + Div_t}{Price_{t-1}} - 1, \quad (9.1)$$

Date	Microsoft (MSFT)			Ferrari (RACE)		
	Price per Share	Dividend per Share	Total Return (%)	Price per Share	Dividend per Share	Total Return (%)
20170131	64.65		4.04	62.13		6.86
20170228	63.98	0.39	-0.43	65.06		4.72
20170331	65.86		2.94	74.36		14.29
20170428	68.46		3.95	75.20	0.68	2.05
⋮	⋮	⋮	⋮	⋮	⋮	⋮
20211029	331.62		17.63	237.17		13.41
20211130	330.59	0.62	-0.12	260.46		9.82
20211231	336.32		1.73	258.82		-0.63

Table 1: Microsoft and Ferrari Stock Return Data

where $Price_{t-1}$ and $Price_t$ are the share prices at the beginning and end of the period, and Div_t , is the dividend, if any, paid during the period.¹ Let's compute the return to a portfolio consisting of both Microsoft and Ferrari during November 2021. To add another practical aspect to the problem, let's assume our portfolio has five shares of Microsoft and ten shares of Ferrari at the beginning and end of November. In other words, we don't buy or sell any shares during the period.

The price of our portfolio at the end of October 2021 (i.e., the start of the period) is

$$Price_{t-1} = \overbrace{5}^{\text{MSFT shares}} \times \underbrace{331.62}_{\text{MSFT price per share}} + \overbrace{10}^{\text{RACE shares}} \times \underbrace{237.17}_{\text{RACE price per share}} = \$4,029.80.$$

The price at the end of November (i.e., the end of the period) is

$$Price_t = 5 \times 330.59 + 10 \times 260.46 = \$4,257.55.$$

While the number of shares of Microsoft and Ferrari in our portfolio didn't change during November, the prices did and, therefore, the portfolio price changed.

Many people will refer to the *value* of their portfolio. But, remember, price and value are only equal if markets are, loosely speaking, efficient. That is, price equals the present value of future cash flows when the market doesn't make a mistake in estimating those cash

¹Strictly speaking, we're assuming the dividend is paid exactly at the end of the period - the last day of the month in this example. This assumption allows us to add $Price_t$ and Div_t without adjusting the dividend for the time value of money. When the time period is sufficiently short, this assumption is of little consequence.

flows or discount rates. Of course, this is rarely true for every firm at every point in time. However, on average, the market does a pretty good job of pricing assets at their true value as evidenced by the difficulty with which investors have in “beating the market.” So, using value and price interchangeably is often not a terrible sin, but it’s important to keep the distinction between the two clear.

Changes in share prices isn’t the only thing that happened to our portfolio in November. Microsoft paid a dividend of \$0.62 per share. Multiplying this cash flow by the number of Microsoft shares equals the total cash flow we received from the company during the period.

$$CashFlow_t = 5 \times 0.62 = \$3.10$$

The total realized return on our portfolio for November 2021 was

$$r_{t-1,t} = \frac{Price_t + CashFlow_t}{Price_{t-1}} - 1 = \frac{4,257.55 + 3.10}{4,029.80} - 1 = 0.0573,$$

or 5.73%.

9.1.2 Portfolio Weights

An equivalent and more convenient way to compute portfolio returns recognizes the return on a portfolio as a weighted sum of the returns to each asset in the portfolio. The weights are the relative amount invested in each asset.

Value-Weighted Returns

Using our example, we have \$4,029.80 invested in our portfolio of Microsoft and Ferrari at the end of October. The fraction invested in each stock is

$$w_{t-1}^{MSFT} = \frac{5 \times 331.62}{5 \times 331.62 + 10 \times 237.17} = \frac{1,658.10}{4,029.80} = 0.4115, \text{ and}$$

$$w_{t-1}^{RACE} = \frac{10 \times 237.17}{5 \times 331.62 + 10 \times 237.17} = \frac{2,371.70}{4,029.80} = 0.5885.$$

The notation w_{t-1}^{MSFT} and w_{t-1}^{RACE} denote the **portfolio weights** as of period $t - 1$. Each weight represent the fraction of the total portfolio value allocated to each asset so the weights sum to one. This is true of every portfolio, no matter how many assets are in the portfolio.

Portfolio weights always sum to one!

The portfolio return is the sum of the portfolio-weighted returns to each asset.

$$\begin{aligned} r_{t-1,t} &= w_{t-1}^{MSFT} \times r_{t,t-1}^{MSFT} + w_{t-1}^{RACE} \times r_{t,t-1}^{RACE} \\ &= 0.4115 \times (-0.0012) + 0.5885 \times 0.0982 \\ &= 0.0573 \end{aligned}$$

We get the same portfolio return, 5.73%, as we did above.

More generally, the realized return on a portfolio, $r_{t-1,t}^P$, containing N (e.g., 2, 10, 100, 10,000) assets during the period $t - 1$ to t can be computed as follows.

$$r_{t-1,t}^P = w_{t-1}^1 r_{t-1,t}^1 + w_{t-1}^2 r_{t-1,t}^2 + \dots + w_{t-1}^N r_{t-1,t}^N \quad (9.2)$$

The portfolio weights, $w_{t-1}^1, \dots, w_{t-1}^N$, represent the fraction of the portfolio value invested in each of the portfolio's N assets as of the start of the period, $t - 1$. The individual asset returns, $r_{t-1,t}^1, \dots, r_{t-1,t}^N$, are the realized returns over the period $t - 1$ to t .

Portfolio Rebalancing

Even if we don't buy or sell any shares, our **portfolio allocation** will change. That is, the weights or relative amount we have invested in each asset in our portfolio will change. Thus, it's important to keep an eye on our portfolio over time to ensure that its composition doesn't deviate too much from what we want. If it does, we need to **rebalance** our portfolio.

Consider our Microsoft and Ferrari portfolio. At the end of October, we had 41.15% of our money invested in Microsoft and 58.85% invested in Ferrari. At the end of November, the value of our Ferrari position will be

$$\underbrace{10 \times 260.46}_{\text{Shares} \times \text{Price}_{\text{Nov}}} = \underbrace{2,371.70 \times (1 + 0.0982)}_{\text{Value} \times (1 + \text{total return})} = \$2,604.40.$$

The value of our Microsoft position will depend on whether and how we reinvest our Microsoft dividend. If we don't reinvest it in the portfolio, the value of our Microsoft position is

$$\underbrace{5 \times 330.59}_{\text{Shares} \times \text{Price}_{\text{Nov}}} = \underbrace{1,658.10 \times (1 - 0.0031)}_{\text{Value} \times (1 + \text{price return})} = \$1,652.95.$$

Notice that we use the *price* return because the dividend was not reinvested.

Alternatively, we can reinvest the \$0.62 per share, $0.62 \times 5 = \$3.10$ in total, to buy $3.10/330.59 = 0.0094$ additional shares of Microsoft. The value of our Microsoft position will be

$$\underbrace{5.0094 \times 330.59}_{\text{Shares} \times \text{Price}} = \underbrace{1,658.10 \times (1 - 0.0012)}_{\text{Value} \times (1 + \text{total return})} = \$1,656.05.$$

Our new portfolio weights are therefore

$$\begin{aligned} w_{t-1}^{MSFT} &= \frac{1,652.95}{1,652.95 + 2,604.60} = 0.3882, \text{ and} \\ w_{t-1}^{RACE} &= \frac{2,604.60}{1,652.95 + 2,604.60} = 0.6118. \end{aligned}$$

Relative to the start of the month, our portfolio is tilted towards or overweighted in Ferrari stock because Ferrari performed well and Microsoft poorly in November of 2022. This is a more general phenomenon. Assets that perform relatively well receive greater weight in our portfolios over time and assets that perform relatively poorly receive less. Depending on our desired allocation, we may need to rebalance our portfolio over time, though any transaction costs or taxes associated with buying and selling stock can have a significant effect on the frequency with which we choose to rebalance.

Of course, we don't have to reinvest the dividend entirely in Microsoft. We could reinvest all of it in Ferrari. Or, we could reinvest it in both Ferrari and Microsoft. The choice is ours. However, regardless of whether or not we reinvest the dividend, we will have to pay taxes on the dividend and it will likely affect our portfolio allocation.

Changing the Portfolio Weights

What happens to our portfolio return as we change the portfolio weights? In other words, if we change how much money we allocate to Microsoft and Ferrari, how does that affect our portfolio return? Figure 9.1 presents the different portfolio returns as we change Microsoft's portfolio weight (w_{t-1}^{MSFT}), i.e., the fraction of our investment allocated to Microsoft as opposed to Ferrari. (The portfolio weight for Ferrari is one minus the Microsoft weight because the weights must always sum to one.)

Specifically, the figure plots the portfolio return, $r_{t-1,t}^P$ as a function of the Microsoft

portfolio weight, w_{t-1}^{MSFT} .

$$\begin{aligned}
 r_{t-1,t}^P &= w_{t-1}^{MSFT} r_{t-1,t}^{MSFT} + w_{t-1}^{RACE} r_{t-1,t}^{RACE} \\
 &= w_{t-1}^{MSFT} r_{t-1,t}^{MSFT} + \underbrace{(1 - w_{t-1}^{MSFT})}_{w_{t-1}^{RACE}} r_{t-1,t}^{RACE} \\
 &= \underbrace{r_{t-1,t}^{RACE}}_{y\text{-intercept}} + \underbrace{(r_{t,t-1}^{MSFT} - r_{t,t-1}^{RACE})}_{\text{Slope}} w_{t-1}^{MSFT},
 \end{aligned}$$

This expression is just the equation of a line, as suggested by the y -intercept and slope annotations.




Figure 9.1: Portfolio Returns with Different Portfolio Weights

The y -intercept corresponds to the portfolio with no money invested in Microsoft ($w_{t-1}^{MSFT} = 0$) and everything invested in Ferrari ($w_{t-1}^{RACE} = 1$). In this case, the portfolio return equals Ferrari's return, 9.82%. Similarly, the x -intercept corresponds to the portfolio with no money invested in Ferrari ($w_{t-1}^{RACE} = 0$) and everything invested in Microsoft ($w_{t-1}^{MSFT} = 1$). That portfolio's return equals Microsoft's return, -0.12%. All points between these two extremes correspond to returns to portfolios containing both Microsoft and Ferrari but in different proportions. For example, the portfolio consisting of equal amounts of Microsoft and Ferrari (i.e., weights of 50%) has a 4.85% return.

Looking at the figure, it's clear that for the month of November we want a portfolio with as little Microsoft as possible. The portfolio return is decreasing in the Microsoft weight (negative slope). Unfortunately, we can't know this future performance for certain, but if

we could or had a strong investment thesis that Ferrari would outperform Microsoft then we should tilt our portfolio accordingly.

9.1.3 Short Positions

What about the points where the weights on Microsoft are negative or greater than 100%? These portfolios contain a **short position** in one of the assets. A short position corresponds to having short-sold an asset. (Recall the discussion from the bond chapter (7)). When we short-sell an asset, we receive money today by borrowing and then selling an asset. But, we have to return that asset in the future. And, depending on whether the price of the asset has gone up or down since we borrowed it, we may have to pay more or less than the value of what we borrowed.

For example, imagine we had \$1,000 at the end of October and we were incredibly **bullish** (i.e., positive, optimistic, expecting to appreciate in value) on Ferrari and **bearish** (i.e., negative, pessimistic, expecting to depreciate in value) on Microsoft. We could short-sell Microsoft stock, essentially betting that the price will fall, and use our \$1,000 *plus* the proceeds from our short-sale to buy Ferrari stock. Let's imagine we short-sell \$200 of Microsoft stock so we can buy \$1,200 of Ferrari stock, and assume that buying and shorting fractional shares is possible.² Our portfolio return for November will be

$$\begin{aligned}
 r_{t-1,t} &= w_{t-1}^{MSFT} \times r_{t,t-1}^{MSFT} + w_{t-1}^{RACE} \times r_{t,t-1}^{RACE} \\
 &= \frac{-200}{1,000} \times (-0.0012) + \frac{1,200}{1,000} \times 0.0982 \\
 &= -0.20 \times (-0.0012) + 1.2 \times 0.0982 \\
 &= 0.1181,
 \end{aligned}$$

or 11.81%. If we short-sell \$200 worth of Ferrari to buy \$1,200 worth of Microsoft, our portfolio return is:

$$\begin{aligned}
 r_{t-1,t} &= w_{t-1}^{MSFT} \times r_{t,t-1}^{MSFT} + w_{t-1}^{RACE} \times r_{t,t-1}^{RACE} \\
 &= \frac{1,200}{1,000} \times (-0.0012) - \frac{200}{1,000} \times 0.0982 \\
 &= 1.20 \times (-0.0012) - 0.2 \times 0.0982 \\
 &= -0.0211,
 \end{aligned}$$

²Buying \$1,200 of Ferrari stock at the end of October means buying $1,200/237.17 = 5.0597$ shares. As of 2023, many trading platforms allow investors to buy fractional shares - e.g., Charles Schwab, Vanguard, Robinhood

or -2.11%. Again, the different returns simply reflect the relative performance of the two stocks in November.

A couple of comments:

- Even with short-selling, the portfolio weights must always sum to one.
- Short-selling an asset is like taking out a loan. The key difference is that we typically know the interest rate on a loan. When we short-sell, we don't know the interest rate because the value of the asset can go up or down and by an arbitrary amount.
- When we borrow money to invest, we are taking a **levered position**, or employing **leverage** or **gearing**. Leverage allows us to increase or **juice** our returns. Without leverage, the highest return we could have earned on our Microsoft-Ferrari portfolio is 9.82% had we put all our money in Ferrari. But, by short-selling Microsoft and taking a levered position in Ferrari, we would have been able to earn more than 9.82%. And, the more leverage we use, the higher that return. The flip side is had we instead short-sold Ferrari to take a levered position in Microsoft, things would have turned out really badly, worse than had we just invested all our money in Microsoft. So, as we'll see more clearly below, leverage increases expected returns but at the expense of additional risk.

9.1.4 Common Portfolio Weights

When people talk about the return on the S&P 500 index, they are talking about the return on a portfolio consisting of the 500 largest U.S. stocks. But, our discussion above begs the question, what are the weights used to construct the S&P 500 index? The answer is: **value weights**. The S&P 500 index is a **value-weighted** portfolio of the 500 largest U.S. stocks. Value weighted portfolios, which are quite common, use the market capitalization of each company in the portfolio scaled by the sum of those market capitalizations.

Consider our Microsoft-Ferrari portfolio. Recall that the market capitalization of a company is the price per share times the number of shares outstanding. As of the end of October 2021, Microsoft's market cap was $331.62 \times 7,507,980,000 = \$2,489,796,327,600$. Ferrari's market cap was $237.17 \times 184,457,000 = \$43,747,666,690$. A value-weighted portfolio of Microsoft and Ferrari as of the end of October would consist of the following portfolio weights.

$$w_{t-1}^{MSFT} = \frac{2,489,796,327,600}{2,489,796,327,600 + 43,747,666,690} = 0.9827$$

$$w_{t-1}^{RACE} = \frac{43,747,666,690}{2,489,796,327,600 + 43,747,666,690} = 0.0173$$

In other words, a value-weighted portfolio of Microsoft and Ferrari means we put 98.3% of our money in Microsoft and 1.7% in Ferrari.

That's an incredibly lopsided allocation. Why? Because Microsoft is much larger than Ferrari - 56.9 times larger! This is what value-weighted portfolios do. They place more weight, i.e., more of your money, in larger stocks. The S&P 500 index is just a bigger version of our Microsoft-Ferrari portfolio. Most of the weight in the S&P 500 is on the very largest stocks. In fact, as of February 2022, the four largest stocks in the index were Apple, Microsoft, Amazon, and Alphabet (a.k.a., Google). These four stocks were responsible for over 20% of the total market cap of the entire index! Most stocks in the index are responsible for a tiny fraction of the index's total market cap, less than a quarter percent to be precise.

So, when we invest in mutual funds or exchange-traded funds (ETFs) that **track** or mimic the S&P 500, we should know that the performance of our investment (i.e., return) is driven largely by the returns of only the largest stocks since most of the money is invested in those large stocks.

Equal-weighted portfolios have portfolio weights that are all the same. An equal-weighted portfolio of Microsoft and Ferrari would invest the same amount of money in both stocks - the weights are 0.50 and 0.50. Likewise, an equal-weighted portfolio of 500 stocks would invest $1 \div 500 = 0.002$ in each of the 500 stocks. Equal weighted portfolios give us equal exposure to each of the stocks in our portfolio. As such, smaller stocks will have the same influence on our portfolio return as larger stocks.

9.1.5 Portfolio Rebalancing

Regardless of whether we are investing our money in a value-weighted or equal-weighted portfolio, the weights of our portfolio will change over time with changes in the valuations of each company. Consider our value-weighted portfolio of Microsoft and Ferrari. At the end of October 2021, the weights on Microsoft and Ferrari are 98.3% and 1.7%, respectively. At the end of November 2021, the weights are 98.1% and 1.9%. This is a small change but a change nonetheless.

Similarly, had we invested \$500 in both Microsoft and Ferrari in October 2021 - an equal-weighted portfolio - then the value of each investment as of November 2021 would be $500 \times (1 - 0.0012) = \499.4 for Microsoft and $500 \times (1 + 0.0982) = \549.1 . The weights for our portfolio at the end of November would be $499.4 / (499.4 + 549.1) = 0.4763$ and $549.1 / (499.4 + 549.1) = 0.5237$ for Microsoft and Ferrari, respectively. The portfolio is no longer equally weighted.

A consequence of these changing weights is that we may want to **rebalance** our portfolio over time. Rebalancing refers to buying or selling different stocks in our portfolio to achieve a desired portfolio composition (i.e., portfolio weights). In our equal-weighted example, we would need to sell some Ferrari stock and use the proceeds to buy Microsoft stock until the weights are once again 50-50, assuming that is our desired portfolio composition.

More generally, if we don't rebalance our portfolio, it will tend to put more weight on those stocks that have done well historically and less weight on those stocks that have done poorly. This might sound reasonable: Why invest in stocks that have done poorly in the past? Well, past returns are not good predictors of future returns; stocks that have done poorly in the past are not more likely to do poorly in the future and similarly for stocks that have done well in the past. Additionally, we'll see that having too much of our savings invested in too few assets is not a good thing because it unnecessarily exposes us to additional risk.

9.1.6 Expected Returns

We saw several different ways to estimate the expected return for an individual stock in chapter 8. To estimate the expected return for a portfolio of stocks - or any collection of assets - we use the same relation that we used for realized returns in equation 9.2 substituting in expected returns for realized returns. Mathematically,

$$r_{t,t-1}^P = w_{t-1}^1 r_{t-1,t}^1 + w_{t-1}^2 r_{t-1,t}^2 + \dots + w_{t-1}^N r_{t-1,t}^N \quad (9.3)$$

where now the returns $r_{t-1,t}^1, \dots, r_{t-1,t}^N$ are expected returns instead of realized returns. Let's use our Microsoft and Ferrari portfolio to illustrate the concept.

The expected return to this two-stock portfolio is

$$r^P = w^{MSFT} r^{MSFT} + w^{RACE} r^{RACE},$$

where the returns all correspond to expected instead of realized returns. The average monthly returns to Microsoft and Ferrari between 2017 and 2021 are 3.11% and 2.87%, respectively. Using these averages as estimates of the expected returns, we can estimate the expected return to a portfolio of Microsoft and Ferrari as

$$r^P = w^{MSFT} 3.11\% + w^{RACE} 2.87\%.$$

What exactly our portfolio expected return is depends on which portfolio we are holding, i.e., which weights, w^{MSFT} and w^{RACE} . Figure 9.2 plots the expected portfolio returns for

all the portfolios in which our Microsoft investment varies between -20% (short Microsoft) and 120% (levered long Microsoft). The two points on the line correspond to 0% weight on Microsoft (i.e., no Microsoft and only Ferrari) and 100% holdings of Microsoft (and no Ferrari). In these cases, the portfolio expected return is exactly equal to the expected return of the one stock in the portfolio.




Figure 9.2: Microsoft-Ferrari Portfolio Expected Returns

It's useful to compare figures 9.1 and 9.2 to reinforce the difference between realized and expected returns. Figure 9.1 is downward sloping because as we have more money invested in Microsoft at the end of October, our portfolio's realized return for November is worse. This result is due to Microsoft's poor performance in November 2021 (-0.12%) compared to Ferrari's strong performance in November 2021 (9.82%). In other words, Figure 9.1 is showing us how different portfolios of Microsoft and Ferrari performed - past tense - in November 2021.

Figure 9.2 is showing us how different portfolios of Microsoft and Ferrari are *expected* - in the future - to perform based on our estimates of their individual expected returns. In this example and throughout this chapter, estimates of individuals expected returns will be based on historical average returns. This choice of estimator is one of convenience. The historical average is easy to understand and calculate, though not necessarily the most accurate. We'll see alternative approaches to estimate expected returns in (chapter 10).

To bring this discussion back to our fundamental value relation, remember that valuation requires discounting *future* cash flows. For that, we need *expected* returns. If we want to

know how an investment has performed, we would look at *realized* returns.

9.2 Portfolio Risk

A portfolio's risk, as measured by its return standard deviation (or volatility), is a little more involved than just taking a weighted average of the individual stocks' standard deviations. It's expression can *look* scary, but it's still just arithmetic and a square root. Let's start with a two-asset example using Microsoft and Ferrari. The volatility of a Microsoft and Ferrari portfolio can be expressed as follows.

$$\sqrt{(w^{MSFT}SD^{MSFT})^2 + (w^{RACE}SD^{RACE})^2 + 2w^{MSFT}w^{RACE}Cov^{MSFT,RACE}} \quad (9.4)$$

As before, the portfolio weights are w^{MSFT} and w^{RACE} . The individual stock volatilities are SD^{MSFT} and SD^{RACE} , where SD denotes standard deviation.

The new term is $Cov^{MSFT,RACE}$, which represents the **covariance** of Microsoft's and Ferrari's returns. Intuitively, covariance measures the relation between two random variables or whether and to what degree they tend to move together. If when Microsoft's return is above its average Ferrari's return tends to be above its average, then the covariance would be positive. If when Microsoft's return is above its average Ferrari's return tends to be below its average, then the covariance would be negative.

Unfortunately, covariance is difficult to interpret because the units are the units of the random variables squared. So, we often rely on **correlation** when describing the strength of any *linear* relation between two random variables.

$$Corr^{MSFT,RACE} = \frac{Cov^{MSFT,RACE}}{SD^{MSFT}SD^{RACE}} \quad (9.5)$$

Like covariance, correlation measures how likely one stock's return will be above its average, if the other stock's return is above its average. If it is more likely, then the returns are **positively correlated**. If it is less likely, then they are **negatively correlated**.

Figure 9.3 illustrates six different random samples of data with a different correlation between the two random variables. For our purposes, we can think of each data point representing the returns to two different assets (e.g., two different stocks or two portfolios - one of stocks and one of bonds). The red line in each graph is a **regression line** that approximates the relation between the two variables.

Correlations are *always* between -1 and 1. When the correlation is either -1 or 1, all of the return pairs lie on a line. Knowing one stock's return implies we can determine exactly

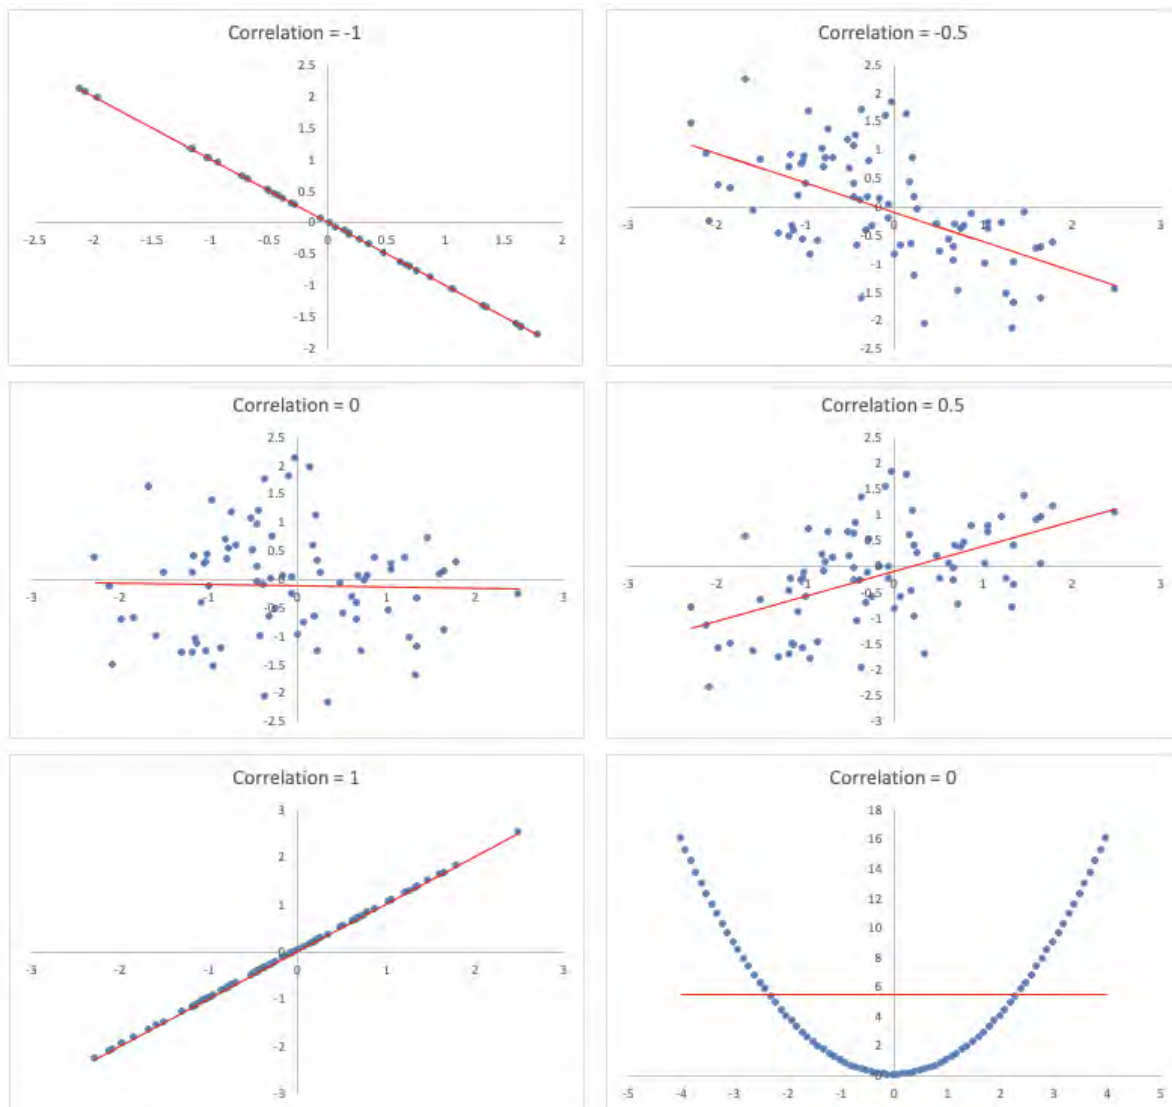


Figure 9.3: Examples of Data with Different Correlations

what the other stock's return is. When two stocks' returns have a correlation of -1 or 1 , the stock returns are said to be **perfectly correlated** or **perfectly linearly related**. As the correlation increases from -1 to 0 , we see the data move away from the line and the slope of the line increase. When the correlation is 0 , then knowing one stock's return provides no information about the other stock's return in a linear sense. In this case, we say the stock returns are **uncorrelated**. As the correlation increases from 0 to 1 , the data moves back towards a line, this time sloping up.

The last figure in the bottom right corner illustrates why we emphasize that correlation measures the *linear* relation between two random variables. In that figure, the y -variable is just equal to the square of the x -variable. These two variables are very closely related, so

much so that knowing x allows us to exactly identify what y must be - x^2 . However, the correlation between the two variables is 0 because of the nonlinear relation. So, correlation measures the strength of the linear relation between two variables but has little to say about variables that may be related in a nonlinear way.

Using equation 9.5, we can express the portfolio volatility in terms of correlation instead of covariance.

$$\sqrt{(w^{MSFT}SD^{MSFT})^2 + (w^{RACE}SD^{RACE})^2 + 2w^{MSFT}w^{RACE}SD^{MSFT}SD^{RACE}Corr^{MSFT,RACE}} \quad (9.6)$$

In the case of Microsoft and Ferrari, the correlation between their returns from 2017 to 2021 is 0.44, suggesting that their returns are moderately positively correlated. Let's use this estimate, and our earlier estimates of volatility to compute the volatility of a portfolio that is has 30% invested in Microsoft and 70% in Ferrari.

$$\sqrt{(0.3 \times 0.0524)^2 + (0.7 \times 0.0780)^2 + 2 \times 0.3 \times 0.7 \times 0.0524 \times 0.0780 \times 0.44} = 0.0631$$

This calculation shows that the monthly volatility of a portfolio consisting of 30% Microsoft and 70% Ferrari is 6.31%.

Figure 9.4 plots the volatility for many Microsoft-Ferrari portfolios. There are several interesting results. First, when our portfolio consists only of Ferrari (Microsoft weight of 0%) or Microsoft (Microsoft weight of 100%), the portfolio volatility equals the volatility of the one stock in the portfolio just as the expected return behaved above. Second, the relation between the portfolio volatility and the portfolio weights is **nonlinear** or curved. Third, this curvature shows that as we load up on either stock, portfolio risk increases. As we start shorting a stock, risk really starts increasing. Remember we discussed leverage and juicing returns above. The curvature of this relation shows the risk implications of leverage. When we are short Microsoft or Ferrari, the risk is relatively high and increasing as we short more of the stock.

Finally, the lowest risk portfolio - the portfolio with the lowest standard deviation - is *not* the portfolio that invests all its money in the low risk stock, i.e., 100% Microsoft. Think about this for a second. This picture shows that we can actually reduce the risk of our investment by adding another stock to our portfolio, even one with a higher standard deviation like Ferrari. The lowest risk portfolio shown in the figure has a standard deviation of 5.07% and invests 81.8% in Microsoft and 100% - 81.8% = 17.2% in Ferrari. Mind blown?

What the heck is going on? **Diversification!!!!** Because the two stocks are only imperfectly correlated (i.e., correlation is not equal to -1 or 1), when Microsoft performs poorly,




Figure 9.4: Microsoft-Ferrari Portfolio Volatility

Ferrari will tend to perform poorly but not all the time (and vice versa). Losses on one stock are occasionally offset by gains on the other thereby reducing the overall risk of the portfolio.

9.2.1 Limits to Diversification

If adding one stock can reduce risk, can adding more reduce risk even further? The answer is yes, assuming the stocks we add aren't too strongly positively correlated? Adding stocks that are perfectly positively correlated - correlation equal 1 - doesn't do anything to help reduce risk because every time one stock is down so are all the others. We need to add stocks, or more generally assets, that are only weakly correlated with one another. Ideally, we'd add stocks that are negatively correlated with one another, but those are rare.

To provide some context, the average correlation between every pair of stocks listed on the New York Stock Exchange (NYSE), National Association of Securities Dealers Automated Quotations (NASDAQ), and American Stock Exchange (AMEX) between 2010 and 2022 is approximately 23.7%, and the average volatility is approximately 40%. Let's look at what happens to the volatility of an equal-weighted portfolio of "average" stocks as we add more and more stocks. In other words, we'll assume each stock has a volatility of 40% and pairwise correlation of 23.7% with every other stock. Figure 9.5 illustrates the results of this exercise.

With just one stock in the portfolio, the volatility of the portfolio equals the volatility




Figure 9.5: The Effects on Portfolio Risk of Adding Stocks Weakly Correlated Stocks

of the one stock - 40%. As we add stocks, the portfolio volatility quickly declines reflecting large diversification benefits initially. With more the 15 stocks the diversification benefits from adding additional stock decreases dramatically. The black dashed line represents the lower bound on the portfolio volatility. In other words, no matter how many stocks we add to the portfolio, we cannot reduce the volatility below that black dashed line. (See the technical appendix for mathematical details.)

This exercise makes two important points. One, we can reduce the risk of our investments by diversifying, that is, adding assets to a portfolio that are not too strongly positively correlated with one another. Two, there are limits to this risk reduction. At a certain point, adding more assets will not reduce risk much further, if at all.

9.2.2 Two Types of Risk: Diversifiable and Non-Diversifiable

Figure 9.5 hints at an important taxonomy of risk that is more clearly illustrated in figure 9.6. There are two types of risk:

1. **Diversifiable risk** is represented by the red area under the curve and above the black dashed line. Diversifiable risk can be eliminated through diversification. More

precisely, diversifiable risk is the volatility that can be reduced by adding assets to our portfolio that are negatively or weakly correlated with our existing assets.

2. **Non-diversifiable risk** is represented by the green area under the black dashed line. Non-diversifiable risk cannot be eliminated through diversification. In other words, non-diversifiable risk is the volatility that we cannot reduce by adding more weakly correlated assets to our portfolio.




Figure 9.6: Diversifiable vs. Non-Diversifiable Risk

Diversifiable risks are those that are uncorrelated across firms. A labor strike, CEO retirement, fire at a warehouse, loss of a customer, supplier bankruptcy, etc. are all examples of risks that companies face. The common feature of these risks is that they won't affect *all* firms at the same time; the risks are uncorrelated with one another. So, when a risk is realized at a company, the stock and bond prices respond because future cash flows are affected. However, other companies are unaffected by the event. Diversifiable risks are also referred to as **unique**, **idiosyncratic**, **firm-specific**, and **non-systematic**. Thus, by diversifying our investments, we can eliminate the effect of these risks on the value of our portfolio. We can reduce the risk of our investments.

In contrast, non-diversifiable risks are correlated across firms. Changes in monetary and fiscal policy, inflation, wars, elections, etc. are risks that affect *all* firms. As such, non-diversifiable risk is often referred to as **systematic** or **market** risk. When inflation strikes,

for example, it affects all companies though to different degrees. Because systematic risk affects all companies at the same time, there is no way to eliminate it through diversification.

Ultimately, diversification is about holding multiple assets that are at weakly correlated with one another to reduce the risk (i.e., volatility) of our investments. Though, there are limits to diversification. Macroeconomic and geopolitical forces create non-diversifiable risks that can't be eliminated just by investing in more assets.

9.3 Risk-Return Tradeoff: Mean-Variance Frontier

Now that we know diversification can reduce risk, there are two questions we need to address. First, is there a cost to diversification? In other words, do we have to sacrifice return for this reduction in risk and, if so, how much? Second, is there an optimal approach to diversifying our investments? In other words, what portfolio of investments should we hold?

9.3.1 The Microsoft and Walmart Frontier

Let's explore the tradeoff between risk and return using a portfolio of Microsoft and Walmart (WMT). Table 2 displays each stock's summary statistics over the period 2012-2016.

	Monthly		Annual	
	MSFT	WMT	MSFT	WMT
Average	1.89%	0.59%	22.73%	7.08%
Standard deviation (SD)	6.45%	4.63%	22.36%	16.05%
Correlation (corr)	(7.33%)		(7.33%)	

Table 2: Microsoft (MSFT) and Walmart (WMT) Return Summary Statistics

The table presents both monthly and annual statistics, the later of which are computed using the scaling results (equation 8.5) from chapter 8. As a reminder, we multiply the average return by 12 and the standard deviation by $\sqrt{12}$ because there are 12 months in a year. The correlation is unaffected the monthly-annual distinction because correlation is unitless. There are a couple of differences relative to the Microsoft-Ferrari portfolio during the 2017 to 2021 period. Microsoft's expected return is lower (1.89% versus 3.11%) and its volatility higher (6.45% versus 5.24%) in the earlier period. Microsoft's correlation with Walmart is significantly lower than that with Ferrari (-0.07 versus 0.44).

Figure 9.7 presents a table and figure. The table presents portfolio weights for Microsoft (w_{MSFT}) and expected returns and volatilities for the corresponding portfolios of Microsoft and Walmart. For example, when we invest 30% of our money in Microsoft - and therefore 70% of our money in Walmart - our expected portfolio return is 11.8% per year give or take 12.7% per year. The figure presents the **mean-variance frontier** for portfolios of Microsoft and Walmart, which is a plot of the expected return-volatility pairs.³ On the horizontal axis is the portfolio standard deviation measuring risk; on the vertical axis the portfolio expected return measuring reward. Each point on the curve corresponds to a different portfolio, several of which are labeled with the portfolio weight on Microsoft. Recall that portfolio weights always sum to one so to determine Walmart's portfolio weight, simply subtract Microsoft's weight from 100%.

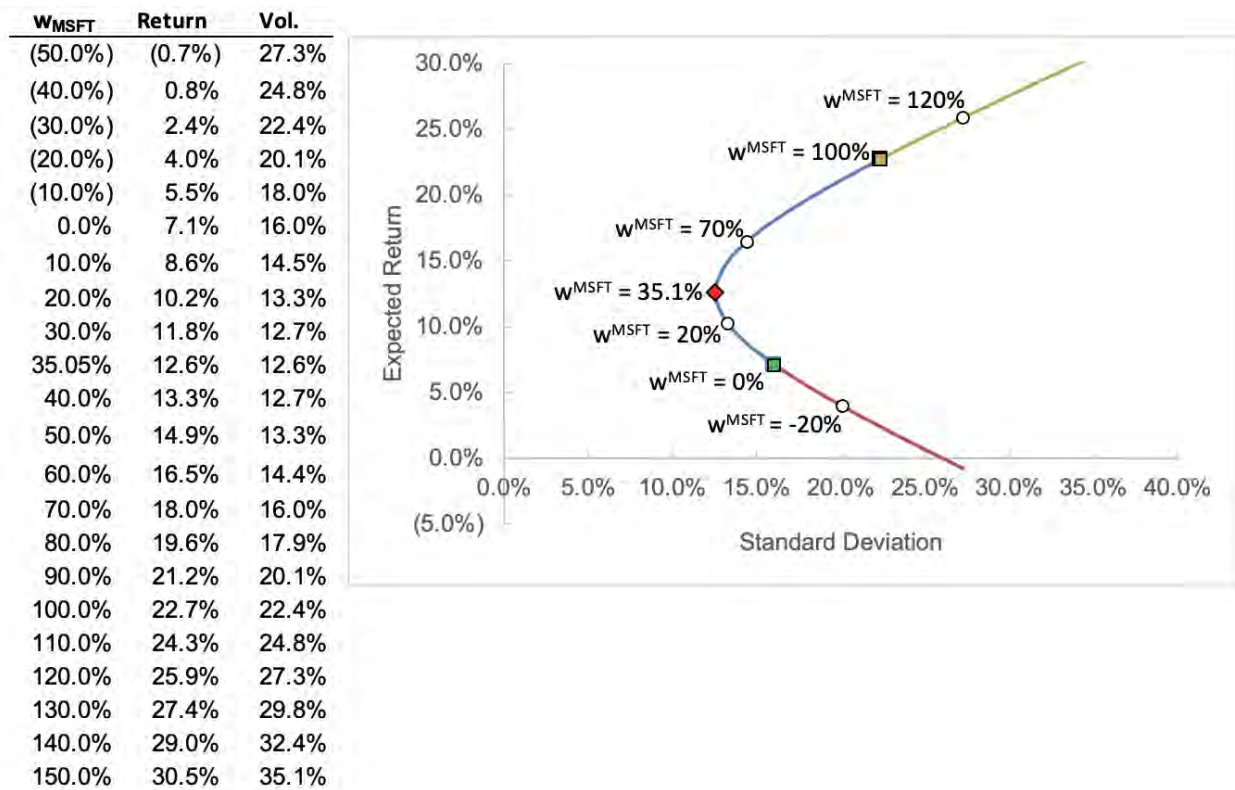


Figure 9.7: Microsoft-Walmart Mean-Variance Frontier

There are three sections of the curve distinguished by different colors.

1. *Short Microsoft, levered long Walmart.* The orange, lower portion of the curve corre-

³This name mean-variance frontier is a bit of a misnomer because the graph uses standard deviation to measure risk, not variance. That said, substituting the variance for standard deviation would only change the units, not the shape of the curve.

sponds to portfolios that are short Microsoft (weight $< 0\%$) and **levered long** Walmart (weight $> 100\%$). The levered long refers to the ownership of Walmart being funded both by our money and borrowed money obtained by shorting Microsoft.

2. *No short positions.* The blue, middle portion of the curve corresponds to portfolios with no short positions (weights on Microsoft and Walmart ≥ 0).
3. *Short Walmart, levered long Microsoft.* The green, upper portion of the curve corresponds to portfolios that are short Walmart (weight $< 0\%$) and levered long Microsoft (weight $> 100\%$).

The green square towards the bottom of the curve corresponds to an investment entirely in Walmart ($w^{MSFT} = 0\%$, $w^{WMT} = 100\%$). The orange square towards the top of the curve corresponds to an investment entirely in Microsoft ($w^{MSFT} = 100\%$, $w^{WMT} = 0\%$). The red diamond corresponds to the portfolio with the lowest standard deviation and is a portfolio we'll discuss in more detail below. Several other randomly chosen portfolios are identified by the white circles.

Most interesting is that the curve is backward bending. As we move along the curve from the bottom right portion to the red diamond something very interesting happens. We are reducing risk (lower standard deviation) and increasing reward (higher expected return)! Double mind blown!!! In other words, if we were just holding Walmart, we can expect to earn 7.08% per year, give or take 16.05%. But, by selling a little bit of Walmart and buying a little bit of Microsoft, we can increase what we expect to earn *and* reduce the volatility of our earnings. For example, a portfolio consisting of 20% Microsoft and 80% Walmart has an annual expected return of 10.2% and volatility of 13.3%.

9.3.2 Efficient and Inefficient Portfolios

This ability to increase expected return and decrease risk suggests another way to look at our portfolios, which is illustrated in Figure 9.8. The figure shows the same curve as in figure 9.7 with only two sections denoted by the red and blue sections of the curve. The bottom, blue section corresponds to the **inefficient** portion of the mean-variance frontier. All the portfolios on this part of the curve are inefficient in the sense that for a given level of risk, there exists another portfolio with the same level of risk but with a higher expected return.

For example, take the portfolio indicated by the point A on the curve. This portfolio has an expected return of 7.1% and a volatility of 16%. If we draw a vertical line from point A




Figure 9.8: Microsoft-Walmart Mean Variance Frontier

back to the curve, we arrive at the portfolio denoted by point B. Portfolio B has the same 16% volatility but an expected return of 18%. Similarly, the portfolio at point C has an expected return of 4% and volatility of 20%. Drawing a vertical line from C back to the curve puts us at point D, which represents a portfolio with the same risk (volatility of 20%) but an expected return of 20.1%. Simply put, no investor should choose a portfolio on the bottom portion of the curve, hence the name *inefficient* portion of the frontier.⁴

The flip side of this discussion is that investors should choose a portfolio on the top part of the curve indicated in red - the **efficient** portion of the frontier. For each of these portfolios, the only way to increase your expected return is by increasing your risk. Take point B. If you want a higher return, you have to travel up and to the right along the curve. But, this means more volatility. Likewise, if you want to reduce your risk from the level at point B, you have to sacrifice expected return. So, for all the portfolios on the top part of the curve, each is *efficient* in that it cannot be improved upon without either increasing risk or sacrificing return.

In general, we like portfolios that are as far Northwest in the figure as possible - portfolios with the highest return and the lowest risk.

⁴There is an exception. **Risk-loving** investors might prefer portfolios A and C to B and D, respectively, because they prefer risk for risk's sake - think gambling addicts. However, most people are **risk-averse**, in which case our conclusion makes sense.

9.3.3 Minimum Variance Portfolio (MVP)

The red diamond in figures 9.7 and 9.8 corresponds to the portfolio with the lowest standard deviation among all possible portfolios. This portfolio is called the **minimum variance portfolio** or **MVP**.⁵ It demarcates the efficient and inefficient frontiers and is an efficient portfolio since there is no other portfolio with the same volatility and higher return.

There are a couple of ways to identify this portfolio - i.e., determine the weights. To avoid unnecessary derivations, let's just focus on the results. For a two asset portfolio - call the assets A and B - the weight on asset A for the minimum variance portfolio is given by the following relation.

$$w^A = \frac{(SD^B)^2 - SD^A SD^B \text{corr}^{A,B}}{(SD^A)^2 + (SD^B)^2 - 2SD^A SD^B \text{Cov}^{A,B}} \quad (9.7)$$

The volatilities of assets A and B are denoted SD^A and SD^B where SD stands for standard deviation. The correlation between the two assets is denoted $\text{Corr}^{A,B}$. The weight on asset B is one minus the weight on asset A, $1 - w^A$, because the weights always sum to one. For completeness, equation 9.7 can be expressed using variances (Var) and covariances (Cov) like so

$$w^A = \frac{Var^B - Cov^{A,B}}{Var^A + Var^B - 2Cov^{A,B}}. \quad (9.8)$$

Let's use equation 9.7 and the data from table 2 to find the MVP for our Microsoft-Walmart portfolios.

$$w^{MSFT} = \frac{0.1605^2 - 0.2236 \times 0.1605 \times 0.0733}{0.2236^2 + 0.1605^2 - 2 \times 0.2236 \times 0.1605 \times 0.0733} = 0.3505$$

The weight on Walmart is therefore $1 - 0.3505 = 0.6495$.

The expected return and volatility of the MVP, like any other portfolio, can be computed using equations 9.3 and 9.4 (or 9.6).⁶ For the Microsoft-Walmart MVP, the expected return and volatility are

$$0.3505 \times 0.2273 + (1 - 0.3505) \times 0.0708 = 0.1257, \text{ and}$$

$$\sqrt{\frac{.2236^2 \times 0.1605^2 - (0.00263)^2}{0.2236^2 + 0.1605^2 - 2 \times 0.00263}} = 0.1257.$$

⁵Since variance is just the square of the standard deviation, the portfolio with the smallest variance is also the portfolio with the smallest standard deviation.

⁶Alternatively, the volatility of the MVP can be computed directly from the variances and covariances of the assets themselves.

$$\sqrt{\frac{Var^A Var^B - (Cov^{A,B})^2}{Var^A + Var^B - 2Cov^{A,B}}}$$

Aside from being the lowest risk portfolio, the MVP as seen above in figures 9.7 and 9.8 distinguishes between efficient and inefficient portfolios. All portfolios whose expected return is equal to or above that of the MVP are efficient. Those with expected returns below the MVP are inefficient.

9.3.4 Changing Correlation

Now let's explore what happens to the mean-variance frontier as the correlation between our two stocks' returns changes. Figure 9.9 shows five different mean variance frontiers for portfolios of Microsoft and Walmart where the portfolios contain *only long positions in at least one stock - no shorting*. (We'll come back to shorting in a minute.) What distinguishes each frontier is the correlation between the two stock's returns, which we've assumed varies from -1.0 to 1.0 by increments of 0.5. The red diamond corresponds to the MVP for each frontier.

Let's start at the extremes - perfect positive (1.0) and negative (-1.0) correlation. In these cases, the stock returns can be written as an **affine** function of the other stock return.⁷

$$ret^{MSFT} = a + b \times ret^{WMT}$$

This equation says that if we know either Microsoft or Walmart's return, we can compute the others stock return using this equation, where a and b are just constants, i.e., numbers we know or can ascertain from observing returns. When the correlation is -1.0, the slope (b) is negative; when 1.0, the slope is positive.

When the correlation is 1.0, the frontier is a line between the two extreme points of 100% Microsoft and 100% Walmart. There is no "backward bend" in the frontier, which is to say there is no benefit to diversification when the correlation is 1.0 *and* we can't take a short position in one of the stocks. Why? Because introducing a perfectly positively correlated stock offers no hedge. When the correlation is 1.0, Microsoft doing poorly means Walmart is doing poorly, and similarly when one stock does well...always. There are no times when one stock does poorly and the other does well precisely because the two stocks are perfectly positively correlated.

In contrast, when the correlation is -1.0, each stock is a perfect hedge for the other. When one stock does poorly, the other stock does well and vice versa. We might think: "If these

⁷An affine function looks like a linear function but, technically speaking, a linear function cannot have non-zero intercept. Further, perfect correlation implies that the random variables are almost surely linearly related - a statistical concept beyond the scope of this book.

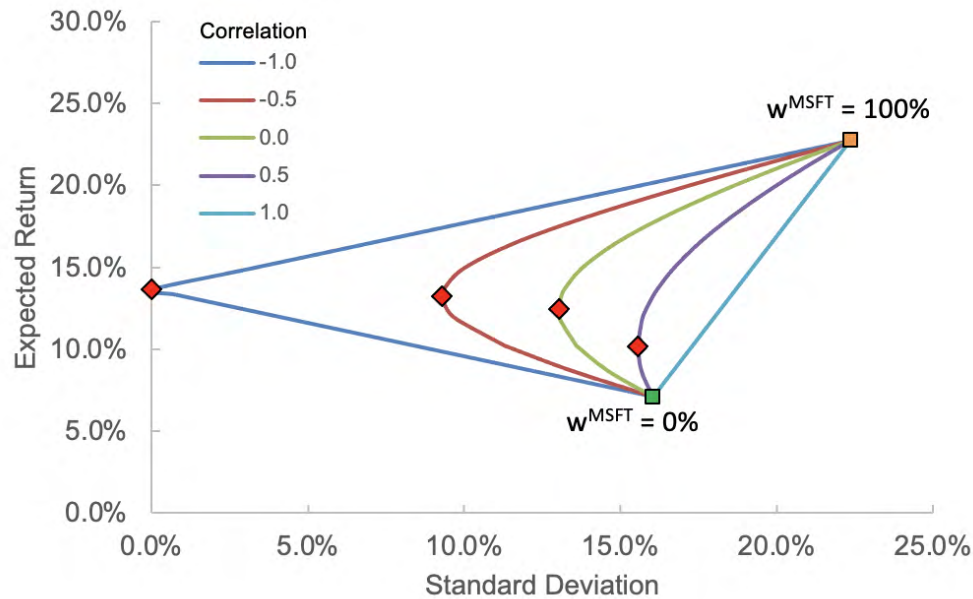


Figure 9.9: Microsoft-Walmart Mean Variance Frontiers with Different Correlations

stocks always move in the opposite direction, how can we ever make money?!?!” The key is to recognize that both stocks have a positive expected return - on average each makes money over time. What the perfect negative correlation tells us is that when one stock has a return above its expected value, the other has a return below its expected value. But, because both of those expected values are positive, we will earn a positive return on average.

Now, perfectly correlated stocks don’t exist in practice. However, the point of the discussion highlights from where the benefits of diversification come. They come from stocks that are not too strongly positively correlated as seen from the other curves. As the correlation between the two stocks’ returns decreases the following equivalent changes to the frontier occur:

- the backward bend in the frontier increases,
- the benefits to diversification increase, and
- the volatility of the minimum variance portfolio decreases.

Notice that there is no MVP indicated on the frontier when the correlation equals 1.0. The MVP when the correlation is high enough can only be achieved by shorting the high volatility stock. For example, when the correlation is 1.0, we have to short more than 254% of Microsoft stock (and therefore buy more than 354% of Walmart - so a heavily levered position) to drive

the portfolio volatility down to zero when the correlation is 1.0. Unfortunately, our expected return for this zero risk portfolio is -32.72%. Not terribly attractive.

In fact, as long as the correlation is 0.72 or larger, we will have to short Microsoft to minimize the volatility of our portfolio. That is, the MVP requires taking a short position in Microsoft as long as the correlation between Microsoft and Walmart is 0.72 or larger.⁸

9.4 A Frontier with Many Assets

An obvious question to ask now is: What happens to the frontier when we have more than two assets in our portfolio? Fortunately, not much. Most of the intuition from our discussion of two asset portfolios carries over to when we have 10, 1,000 or 10,000 assets in our portfolio. Additionally, it's important to emphasize that while our discussion has and will continue to focus on stocks, there is nothing in what we are saying that is unique or specific to stocks. We could have a portfolio of stocks, bonds, real estate, cryptocurrency, venture capital, private equity, art, wine, etc. The analysis and intuition is unchanged.

Let's construct a mean variance frontier from the following 12 stocks to illustrate the similarities and differences with the two asset example on which we've focused.

- Microsoft (MSFT)
- International Business Machines (IBM)
- General Mills (GIS)
- Caterpillar (CAT)
- Boeing (BA)
- Archer Daniels Midland (ADM)
- Walmart (WMT)
- Hershey (HSY)
- Proctor & Gamble (PG)
- Deere & Co (DE)
- JP Morgan Chase (JPM)
- EBAY (EBAY)

We'll use monthly data from January 2012 to December 2016 to estimate expected returns, standard deviations, and correlations. The math for constructing the frontier with more than two assets can be found in the technical appendix for those interested. Our discussion here focuses on the results and intuition.

Figure 9.10 presents the frontier and some key data points. The blue circles correspond to the average return-volatility pairs for the 12 stocks. Each data point is labeled with the stock's ticker symbol. For example, Hershey's (HSY) average annual return and volatility are 15.35% and 18.10%, respectively. The green circle labeled EWRET represents the average

⁸How we got 0.72 is somewhat technical and left to the accompanying chapter Excel spreadsheet, which shows how to compute this number.

return and volatility of an equal-weighted portfolio of the 12 stocks. In other words, the 12 portfolio weights all equal to $1/12=0.0833$.




Figure 9.10: Mean Variance Frontier with Multiple Assets

The red and blue curve is the mean variance frontier for the 12 stocks. Two portfolios on the frontier are labeled “Eff 1” and “MVP”. The first is an arbitrarily chosen efficient portfolio, hence the label Eff. The second portfolio is the minimum variance portfolio. The portfolio weights for these two portfolios are presented in Table 3. The table shows that the portfolios are quite different in that their weights on the different stocks are quite different, despite being relatively close to one another on the frontier. Also notice that the weights of the two portfolios both sum to one.

The equal-weighted portfolio of the 12 stocks (EWRET) is *not* on the frontier, though it is closer than any individual stock. But, what can we really say about EWRET relative to the individual stocks? The equal-weighted portfolio is unambiguously better than GIS, ADM, DE, PG, WMT, IBM, and CAT. Why? Because EWRET has a higher expected return than all of these stocks *and* lower volatility. However, it is not necessarily preferable to HSY, BA, MSFT, JPM, and EBAY. Why? Because EWRET may have lower volatility than these stocks but it also has a lower expected return.

Stepping back, the frontier for our 12 assets looks a lot like our frontier for two assets (9.8). The key difference is that the two-asset frontier represents *all* of the portfolios that can be formed with those assets. Every portfolio of a two-asset portfolio lies on the frontier -

Stock Ticker	Eff 1	MVP
MSFT	12.86%	8.79%
ADM	0.68%	-0.23%
IBM	-5.88%	8.12%
HSY	10.50%	9.50%
GIS	17.82%	14.09%
PG	7.79%	15.56%
CAT	-1.47%	0.72%
DE	14.73%	18.19%
BA	8.07%	6.46%
JPM	3.55%	-5.68%
WMT	21.44%	18.46%
EBAY	9.91%	6.01%

Table 3: Efficient Portfolios (Negative values in parentheses)

some on the efficient part, some on the inefficient part. The frontier for more than two assets contains a *subset* of portfolios that can be formed. The remainder of the portfolios fall in the interior of the frontier indicated by the orange shaded area. In other words, our **investment opportunity set**, which consists of all the portfolios that can be formed, consists of two parts:

1. the portfolios *on* the frontier (red and blue curve), and
2. the portfolios *within* the frontier (orange shading).

As investors, the portfolios in which we should be interested are those falling on the efficient frontier indicated by the red curve (portfolios with expected returns greater than or equal to that of the MVP). These efficient portfolios offer the highest expected return for a given level of risk or, equivalently, the lowest risk for a given level of expected return.

Notice in figure 9.10 that the benefits to diversification are quite large in the following sense. We can reduce risk substantially without sacrificing return by moving from any individual stock to the efficient frontier. For many stocks - PG, WMT, DE, IBM, CAT - we can not only reduce risk, we can *increase* return by moving to the efficient frontier. This is what diversification is all about, and why people are strongly encouraged to hold diversified portfolios - the ability to significantly reduce risk without sacrificing return and at time increasing it.

9.4.1 A New Measure of Risk - Covariance

Up to now, we've focused on volatility (i.e., standard deviation) as our measure of risk. But, let's consider the volatility of a portfolio as the number of assets increases, and because volatility is just the positive square root of variance, let's focus on variance to avoid having to put equations under the square root symbol.

The variance of a 2-asset portfolio (e.g., equation 9.4) can be expressed as the sum of the two weighted variances plus two covariance terms.

$$Var^P = (w^1)^2 Var^1 + (w^2)^2 Var^2 + 2w^1 w^2 Cov^{1,2}$$

The variance of a 3-asset portfolio is the sum of the three weighted variances plus six covariance terms.

$$\begin{aligned} Var^P = & (w^1)^2 Var^1 + (w^2)^2 Var^2 + (w^3)^2 Var^3 \\ & + 2w^1 w^2 Cov^{1,2} + 2w^1 w^3 Cov^{1,3} + 2w^2 w^3 Cov^{2,3} \end{aligned}$$

(Last one) The variance of a 4-asset portfolio is the sum of the four weighted variances plus 12 covariance terms.

$$\begin{aligned} Var^P = & (w^1)^2 Var^1 + (w^2)^2 Var^2 + (w^3)^2 Var^3 + (w^4)^2 Var^4 \\ & + 2w^1 w^2 Cov^{1,2} + 2w^1 w^3 Cov^{1,3} + 2w^1 w^4 Cov^{1,4} \\ & + 2w^2 w^3 Cov^{2,3} + 2w^2 w^4 Cov^{2,4} + 2w^3 w^4 Cov^{3,4} \end{aligned}$$

If we continue this process, we'll see that a portfolio with N assets has a variance consisting of N variance terms and $N^2 - N$ covariance terms.

Every time we add another asset to our portfolio, the number of variance terms increases by one, but the number of covariance terms increases by $2(N-1)$.⁹ For example, the variance of portfolio with 100 assets has 100 variance terms and 9,900 covariance terms. Adding one more asset to the portfolio adds one more variance term and 200 covariance terms! Therefore, for large, well-diversified portfolios, the variance, and therefore volatility, of the portfolio is determined largely by the covariances between all of the assets as opposed to the variances of the assets. There are so many more covariance terms that the volatilities almost cease to matter after a certain point.

⁹There are N variance terms and $N^2 - N$ covariance terms in the variance expression for a portfolio with N assets. The increase in the number of covariance terms each time we add another asset is therefore $N^2 - N - [(N-1)^2 - (N-1)] = 2(N-1)$.

Here's the punchline. If we're considering adding an asset to our portfolio, the volatility of the asset is less relevant than its covariance (or correlation) with our portfolio. In other words, the relevant measure of risk for an asset that we are considering adding it to our portfolio is *not* volatility, it's covariance with our portfolio! We have another measure of asset risk - covariance or correlation. And, to be perfectly clear, let's summarize when to use each.

- We use volatility as a measure of an asset's risk when we are assessing the asset by itself, independent of any other assets.
- We use covariance (or correlation) as a measure of an asset's risk when we are assessing the asset in the context of other assets.

For example, if we want to know how risky our savings are, we want to know the volatility of our savings. If we want to know how risky a new investment is, we want to know how correlated it is with our existing investments. This is important. Once we own assets such as a house, stocks, bonds, etc., the volatility of new investment opportunities becomes less important than the correlation of those opportunities with our existing assets because it is the correlation that will contribute the most to the volatility of our total wealth.

9.5 Adding a Risk-Free Asset to Our Portfolio

Notice in figure 9.10 that there are no portfolios on the vertical axis. The portfolio closest to the vertical axis is the MVP. This means there are no **risk-free** investments, or equivalently investments with zero volatility. We saw earlier that if two stocks were perfectly correlated - correlation of -1.0 or 1.0 - then we could construct a risk-free portfolio. But, in practice, there are no perfectly correlated stocks (or bonds).

The closest thing to a tradeable risk-free asset is a Treasury security issued by the U.S. federal government. Though, a Treasury security is *not* risk-free. Even if the probability of the government defaulting was zero - it is not - Treasuries are subject to interest rate, inflation, and in the case of notes and bonds, reinvestment risk (see chapter 7). Further, this risk can be substantial for longer maturity bonds. So, take the reference to Treasury's as risk-free securities with a large grain of salt.¹⁰

¹⁰Theoretically, a risk free security would be an inflation protected perpetuity with zero probability of default. See Campbell, J. Y., and Viceira, L. M. (2001): Who should buy long-term bonds? American Economic Review 91, 99–127.

Figure 9.11 illustrates what happens to our two-asset portfolio when a risk-free security is added to the investment opportunity set. The risk-free asset is identified by the purple square on the vertical axis implying that the return never varies (grain of salt...). Using a 30-day T-bill, the risk-free return is an estimated 1.3% per annum over the 2012-2016 period.




Figure 9.11: Mean Variance Frontier with Risk-Free Asset

Portfolios consisting of the risk-free asset (T-bill) and Microsoft are represented by the red line, which is called a **capital allocation line** or **CAL**. A CAL describes the risk-return tradeoff for portfolios consisting of a risky asset - possibly a portfolio of risky assets - and a risk-free asset. Let's consider three sections of the CAL in Figure 9.11. The first is the upward sloping red line between the purple and orange squares. This section represents all of the portfolios that are long the risk-free asset, Microsoft, or both. That is, both portfolio weights are greater than or equal to 0).

The section of red line extending above and right of the orange square represents portfolios that are short the risk-free asset and levered long Microsoft - negative portfolio weight on the T-bill, portfolio weight greater than 100% on Microsoft. These portfolios are formed by borrowing money at the risk-free rate, combining this money with any other money, and investing everything in Microsoft. Finally, the red line starting at the purple square and sloping downward represents portfolios that are short Microsoft and levered long the risk-free asset. Hopefully, we can recognize that this last group of portfolios is inefficient because there are portfolios offering higher returns with equal or less risk. Interpretation of the green lines is similar except they represent portfolios of the risk-free asset and Walmart stock.

Notice the following changes when we introduce a risk-free asset. First, we can achieve a risk-free portfolio by putting all our money in the risk-free asset (i.e., Treasury's). Second, our investment opportunity set has grown. With only two assets - Microsoft and Walmart stock - the only portfolios in which we could invest - and therefore the only expected return-standard deviations we could hope to experience - fell on the blue curve. Now we can also invest anywhere along the red or green lines, in addition to the blue curve. Finally, the efficient portion of the frontier is different and worth discussing in some detail.

9.5.1 Temporary New Efficient Frontier

With two stocks, the efficient frontier consisted of all the portfolios with expected returns greater than or equal to that of the minimum variance portfolio (the red curve in figure 9.8). But, consider a portfolio that shorts Walmart to go levered long Microsoft - the blue curve up and to the right of the orange square in figure 9.11. We would never want to be on this part of the curve, because we could achieve a higher expected return for the same risk by shorting the T-bill and going levered long Microsoft (i.e., jumping up to the red line). Similarly, notice that the red line crosses the blue curve just above the MVP. That tiny bit of blue curve between the MVP and this point of intersection is also no longer efficient.

The new efficient frontier consists of three sections:

1. All the portfolios on the red line between the purple square and the blue curve just above the MVP. These are portfolios that are long the T-bill and long Microsoft.
2. All the portfolios on the blue curve between the point where the red line first intersects the blue curve (just above the MVP) and the orange square. These portfolios are long Microsoft and Walmart.
3. All the portfolios on the red line up and to the right of the orange square. These portfolios are short the T-bill and levered long Microsoft.

Figure 9.12 isolates the new efficient frontier with a risk-free asset.

This subsection refers to this efficient frontier as temporary for a reason; we can do better. Let's see how.




Figure 9.12: Efficient Frontier with Risk-Free Asset

9.5.2 Sharpe Ratio

The expected return for our portfolio of a T-bill and Microsoft stock, r^P , is like any other portfolio expected return, a weighed sum of the individual expected returns (equation 9.3).

$$r^P = w^F r^F + w^{MSFT} r^{MSFT}, \quad (9.9)$$

The weight and expected return to the risk-free asset are w^F and r^F , respectively. Theoretically speaking, there is no difference between the expected return and realized return for a risk-free asset because the risk-free asset return never varies - volatility is equal to 0. In other words, the risk-free return is always the same number. Of course, this is a purely theoretical assumption as we noted - grain of salt...

Using equation 9.6, we can compute the volatility of our portfolio, SD^P .

$$\sqrt{SD^P} = \sqrt{(w^F SD^F)^2 + (w^{MSFT} SD^{MSFT})^2 + 2w^{MSFT} w^F SD^{MSFT} SD^F Corr^{MSFT,F}} \quad (9.10)$$

But, remember our definition of a risk-free asset - zero volatility. This assumption means that $SD^F = 0$. It also means that $Corr(MSFT, F) = 0$ because no matter what Microsoft's return does, the risk-free rate is always the same. That is, Microsoft's, and every other asset's, return is uncorrelated with the risk-free return. After substituting zeroes for the risk-free volatility and correlation into equation 9.10, the volatility of our T-bill-Microsoft portfolio is

$$SD^P = \sqrt{(w^{MSFT} SD^{MSFT})^2} = w^{MSFT} SD^{MSFT} \quad (9.11)$$

Now, let's pay close attention to the red line in figure 9.13, which shows the expected return and volatility combinations for portfolios of the risk-free asset and Microsoft. (This is the same red line as in figure 9.11). We can combine equations 9.9 and 9.11 to get an




Figure 9.13: Portfolios of T-bill and Microsoft

equation for this line.

$$\underbrace{r^P}_{y\text{-variable}} = \underbrace{r^F}_{y\text{-intercept}} + \underbrace{\frac{r^{MSFT} - r^F}{SD^{MSFT}}}_{\text{Slope (Sharpe Ratio)}} \times \underbrace{SD^P}_{x\text{-variable}} \quad (9.12)$$

The y -variable is the expected return on the portfolio. The y -intercept equals the risk-free rate, where the line touches the y -axis. The x variable in equation 9.12 is the portfolio volatility.

The slope of our line is the ratio of the change in height to the change in length - remember “rise over run?” The change in height is the expected return on Microsoft minus the risk-free rate. The change in length is the volatility of Microsoft minus zero, or just the volatility of Microsoft. This slope has an interesting interpretation and is also referred to as the **Sharpe Ratio**, named after the Nobel Prize-winning financial economist William Sharpe. The numerator of the Sharpe ratio measures an asset's **risk premium** or **excess expected return**, defined as the expected return minus the risk-free rate. Microsoft's risk premium based on its historical performance is $22.73\% - 1.3\% = 21.43\%$. The denominator measures the volatility or risk of the asset - 22.36% again based on Microsoft's historical

performance. Therefore, the Sharpe Ratio measures the risk premium or excess expected return per unit of risk - 0.96 for Microsoft.

All else equal, higher Sharpe ratios on our investments are better because they offer more return per unit of risk. Looking at figure 9.11, we can see that the red line connecting the risk free-asset and Microsoft is more steeply sloped (i.e., higher Sharpe ratio) than the green line connecting the risk-free asset and Walmart - suggesting that Microsoft has a higher Sharpe ratio. Related, portfolios of T-bills and Microsoft are strictly better than portfolios of T-bills and Walmart because they offer higher returns for the same risk.

9.5.3 Efficient Frontier

Remember, we like portfolios that are as far Northwest as possible in the expected return-volatility plane - portfolios with the highest expected return and the lowest risk. Figure 9.14 illustrates how *maximizing the Sharpe ratio* achieves that goal.




Figure 9.14: Maximizing the Sharpe Ratio

The blue curve is our Microsoft-Walmart frontier that we've seen in previous figures. The lines, or rays, emanating from the risk-free return on the y-axis are mean-variance frontiers or capital allocation lines corresponding to different portfolios of risky assets and the risk-free asset. Take ray A for example. This ray shows the risk-reward combinations (standard deviation and expected return) for portfolios containing a T-bill and Walmart stock, i.e., a portfolio of Microsoft and Walmart stock in which the weights are 0% and 100%, respectively.

The left most point on the ray corresponds to a portfolio that is 100% invested in the T-bill. As we move up and to the right along the ray, we are reallocating our money from the T-bill to Walmart stock - the portfolio weight on the T-bill decreases, the portfolio weight on Walmart stock increases. The point at which the ray intersects the blue curve - the green square - is a portfolio in which 100% of our money is invested in Walmart stock. Beyond the green square, we are shorting the T-bill (i.e., borrowing money) and going levered long Walmart stock.

The slope of ray A corresponds to Walmart's Sharpe ratio equal to Walmart's excess expected return divided by its volatility. Substituting our estimates for Walmart's expected return, the risk-free rate, and Walmart's volatility produces $(7.1\% - 1.3\%)/16.0\% = 0.36$. Investing in Walmart stock generates 0.36 units of excess return per unit of risk.

The other rays are similar except they have a different composition of risky assets. Take ray B. The left most point on the ray corresponds to a portfolio that is 100% invested in the T-bill. As we move up and to the right along the ray, we are shifting our money out of the T-bill and into a portfolio of Microsoft and Walmart. Which portfolio? One that is 10% invested in Microsoft and 90% in Walmart as indicated by the ray's label in the figure. Where ray B crosses the blue curve - the light orange square - is the point at which all of our money is invested in the Microsoft-Walmart portfolio, split 10-90 between the two stocks. Beyond the light orange square we are shorting the T-bill and going levered long in the Microsoft-Walmart portfolio.

We have a portfolio containing a portfolio! What exactly does this mean? Imagine we have \$100 to invest and we are holding a portfolio that is split 50-50 between the T-bill and the Microsoft-Walmart portfolio. In the figure, this point would be on ray B approximately half way between the purple square and the light orange square. Because our money is split 50-50, \$50 is invested in T-bills. The other \$50 is invested in the Microsoft-Walmart portfolio, which allocates 10% to Microsoft ($0.10 \times 50 = \$5$) and 90% to Walmart ($0.90 \times 50 = \45). So, our asset allocation is \$50 in T-bills, \$5 in Microsoft stock, and \$45 in Walmart stock.

The Sharpe ratio for ray B is the excess expected return for the Microsoft-Walmart portfolio divided by the volatility of the Microsoft-Walmart portfolio. The expected return for this portfolio can be found using equation 9.3; the standard deviation using 9.6. Using

the information from table 2 we can estimate these statistics.

$$\begin{aligned}\mathbb{E}(r^P) &= 0.10 \times 22.73\% + 0.90 \times 7.08\% \\ &= 8.65\% \\ SD^P &= \sqrt{(0.10 \times 0.2236)^2 + (0.90 \times 0.1605)^2 + 2(0.10)(0.90)(0.2236)(0.1605)(-0.0733)} \\ &= 14.45\%\end{aligned}$$

The risk-free return is 1.3%. Therefore, the Sharpe ratio for our 10-90 Microsoft-Walmart portfolio is $(8.65\% - 1.3\%)/14.45\% = 0.51$.

Rays C through F are constructed and interpreted similarly, differing only in the portfolio of risky assets - Microsoft and Walmart - that are combined with the T-bill. The message of figure 9.14 is that investors want to be on the ray with the steepest slope, i.e., the highest Sharpe ratio. Why? Pick a point on any ray other than that with the highest Sharpe ratio. For the same risk, we can find a portfolio with a higher expected return by moving up to a more steeply sloped ray. Alternatively, for the same expected return, we can find a portfolio with a lower volatility by moving left to a more steeply sloped ray.

Focusing on the steepest sloped ray, F, reveals that it is **tangent** to the blue curve; it just barely touches the curve at a single point. The point at which it touches the curve, indicated by the turquoise circle, is called rather creatively the **tangency portfolio**. In this example, the tangency portfolio has weights of 62.1% and 37.9% on Microsoft and Walmart. Details of how to find this portfolio are in the technical appendix.

There are no rays more steeply sloped than F because there are no investment opportunities outside of the blue curve. In other words, drawing a ray more steeply sloped than ray F is nonsensical because there are no investments with expected returns and volatility pairs that *exist* outside of the blue curve. We also can't "improve" on any portfolios found on ray F. For any portfolio on ray F, if we want a higher return, we have to accept more risk by moving up and to the right along the ray. If we want less risk we have to accept a lower return by moving down and to the left along the ray.

Importantly, this conclusion holds true when we consider more than two risky assets. Figure 9.15 presents the efficient frontier for the 12 stocks discussed above in section 9.4. The efficient frontier is the red ray emanating from the risk-free return on the vertical axis that is tangent to the frontier of risky assets. This ray contains risk-reward combinations for portfolios of the risk-free asset and the tangency portfolio, whose weights are shown in Table 4.

Points on the ray between the purple square and the turquoise circle are long both risk-free asset and tangency portfolio. Points on the ray up and to the right of the turquoise

circle are short the risk-free asset and levered long the tangency portfolio. The downward sloping green ray emanating from the purple square are short the tangency portfolio and levered long the risk-free asset - not a good strategy because for every portfolio on the green ray there is another portfolio on the red ray that offers the same risk but greater reward.




Figure 9.15: Efficient Frontier for 12 Stocks and Risk-free Asset

9.5.4 Different Lending and Borrowing Rates

In the real world, people borrow and lend at different rates. When you want to borrow money, the lowest borrowing rate you can get is typically higher than the highest lending rate you can get. For example, if we look to borrow money in the form of a mortgage and at the same time look to lend money by investing in a savings account or buying a Treasury security, the interest rate on our mortgage will be higher than the interest rate on our savings account or Treasury.

However, the rays emanating from the risk-free rate in the last several figures assume these rates are the same. This assumption can be seen in equation 9.12, which shows only one risk-free rate, r^F . It can also be seen in the figure because each ray has the same slope at every point. Figure 9.16 shows what happens to these rays when there are different borrowing and lending rates.

Stock	Portfolio Weight (%)
Microsoft (MSFT)	31.24
Archer Daniels Midland (ADM)	4.81
International Business Machines (IBM)	-69.07
Hershey (HSY)	15.013
General Mills (GIS)	34.61
Proctor & Gamble (PG)	-27.31
Caterpillar (CAT)	-11.42
Deere & Co (DE)	-0.84
Boeing (BA)	15.31
JP Morgan Chase (JPM)	45.25
Walmart (WMT)	34.85
EBAY (EBAY)	27.54

Table 4: Tangency Portfolio Weights

The blue curve is our usual risky-asset frontier consisting of Microsoft and Walmart. The tangency portfolio is indicated by the turquoise circle. The purple square on the vertical axis is the risk-free return of 1.3%. However, now we can only lend (i.e., invest) at this rate, hence the notation r^{LEND} . That is, we can buy a Treasury security, but no one is going to lend to us at the same interest rate. If we want to borrow money, we're going to have to pay a higher interest rate, which is indicated in the figure by the green square at let's say 8%, r^{BORROW} label.

The portfolios on the black line between the purple square and the turquoise circle correspond to portfolios in which we are lending (i.e., long the T-bill) and long the tangency portfolio. The portfolios on the black dashed line, which extends the solid black line, correspond to portfolios in which we are borrowing (i.e., short the T-bill) and taking a levered long position in the tangency portfolio. But, shorting the T-bill is the same thing as borrowing money and we can no longer borrow at 1.3%. So, portfolios in which we are short a T-bill at 1.3% are no longer a possibility for us. The black dashed line represents only hypothetical portfolios.

The portfolios on the dashed red line are similarly hypothetical. We can't lend at 8%, i.e., earn a risk-free return of 8%. (Oh, but I wish we could today in late 2021. Update...we're getting close in 2023:)). Therefore, the portfolios between the green square and the turquoise circle - portfolios that are lending at 8% and long the tangency portfolio, are not options for us. The solid red line that extends the dashed red line consists of portfolios that are borrowing




Figure 9.16: Efficient Frontier for 12 Stocks and Risk-free Asset

at 8% and going levered long the tangency portfolio. These portfolios are available to us. However, we would never want to invest in these portfolios because there are portfolios of just the risky assets on the blue curve that dominate the portfolios on the solid red line - offer a higher expected return for the same risk or lower risk for the same expected return.

Mathematically, we have two equations for these solid lines. When we are long the risk-free asset ($w^F \geq 0$), we are lending money and the equation for the solid black line is

$$r^P = r^{LEND} + \frac{r^T - r^{LEND}}{SD^T} SD^P.$$

This equation says the expected return for portfolios consisting of the risk-free asset and the tangency portfolio, T , when we are long the risk-free asset (i.e., buying Treasury securities) equals

1. the return on the risk-free asset r^{LEND} , plus
2. the Sharpe ratio of the Tangency portfolio at the lending rate times the volatility of the overall portfolio - Treasury and Tangency.

When we are short the risk-free asset ($w^F < 0$), we are borrowing money and the equation

for the solid red line is

$$r^P = r^{BORROW} + \frac{r^T - r^{BORROW}}{SD^T} SD^P.$$

This equation says the expected return for portfolios consisting of the risk-free asset and the tangency portfolio, T , *when we are short the risk-free asset* (i.e., borrowing money) equals

1. the cost of borrowing money r^{BORROW} , plus
2. the Sharpe ratio of the Tangency portfolio at the borrowing rate times the volatility of the overall portfolio - Treasury and Tangency.

The main message of this analysis is that in practice it is important to recognize the differences between borrowing and lending rates when we invest. As figure 9.16 shows, these different rates change the efficient frontier. For some institutional investors (e.g., hedge funds), the difference between borrowing and lending costs can be quite small thereby having a small effect on the efficient frontier. For most individuals, the difference tends to be quite large, requiring a moment of reflection before borrowing money to invest in risky assets.

That said, many homeowners carry a mortgage and have savings. This situation is similar to shorting a (near) risk-free asset and taking a levered long position in risky assets. Why do people do this? Many reasons. Some people want to maintain liquidity for unexpected expenses; it can be very difficult and costly to get money out of house. Some people believe that over long horizons they can earn more on their savings than what they pay on their mortgage; and for those with mortgage rates below 3% chances are they're correct. Some may want to avoid too much concentration of their wealth in one asset - their home - for diversification purposes. Finally, some may simply not realize what they're doing.

9.6 What Diversification Does *Not* Mean

9.6.1 Return Risk vs. Dollar Risk

We've shown that by investing in many assets that are not too highly correlated with one another, we can reduce our risk exposure without sacrificing return. This phenomenon is diversification. However, lurking in the background of this discussion is the assumption that *we do not change the size of our investment*.

Take our Microsoft and Walmart example from above. Imagine we have \$100 invested in Microsoft. The annual expected return and volatility of Microsoft was estimated as 22.73%

and 22.36%, respectively. (See table 2.) This means that one year from today we can expect to earn \$22.73, give or take \$22.36.

Diversification suggests that rather than investing \$100 in Microsoft, we should consider investing a portion of that \$100 in Microsoft and the rest in Walmart (or many other assets). Doing so can reduce our risk. What diversification does *not* say is that we should invest *another* \$100 in Walmart, bringing our total investment up to \$200. By increasing our investment, we are in fact increasing our risk in terms of how much money we can lose. Let's illustrate using results from above.

The expected return and volatility of an equal-weighted portfolio of Microsoft and Walmart are 14.9% and 13.3%. In other words, by splitting our money between Microsoft and Walmart, we can reduce the volatility of our investment from 22.36% to 13.3%. If we apply these results to our original \$100 investment, then the dollar volatility of our investment is $0.2236 \times 100 = \$22.36$ if we invest entirely in Microsoft and $0.1333 \times 100 = \$13.33$ if we invest \$50 in Microsoft and \$50 in Walmart.

However, if we invest \$100 in Microsoft and \$100 in Walmart - an equal weighted portfolio that is twice the size of our original investment, then the dollar risk increases even though the return risk does not increase. In other words, the volatility of an equal weighted portfolio 13.33%, regardless of how much we invest. However, the more we invest, the greater the risk in terms of dollars we could lose. Our \$200 portfolio has a dollar volatility of $0.1333 \times 200 = \$26.66$. This amount is greater than that of our \$100 equal-weighted portfolio - no surprise. It is also greater than our \$100 investment entirely in Microsoft (\$22.36).

This is all to say that diversification is about taking a fixed amount of money and spreading it across many weakly correlated assets. It is not about investing more money in many different assets.

9.6.2 Time Diversification

Now, instead of increasing the size of our investment, let's consider the effects of increasing the duration of our investment. More simply, can we reduce risk by investing for longer horizons, so-called **time diversification**? Unfortunately no. Table 5 presents some interesting facts about investing in the Standard and Poor's (S&P) 500 index between 1927 and 2022. The S&P 500 is a value-weighted portfolio of the 500 largest publicly traded stocks and a popular barometer of stock market performance.

The table shows several empirical facts about stock returns as the investment horizon increases. First, expected returns increase, but so too does volatility. In other words,

Horizon (Years)	Average Return	Return Volatility	Probability of Loss	Average loss when loss
1	11.82	19.36	26.60	-13.31
2	24.45	29.70	17.02	-22.30
5	71.86	59.23	11.70	-19.70
10	199.03	139.06	5.32	-10.45
20	815.74	550.78	0.00	N/A

Table 5: S&P 500 Performance Statistics Over Different Investment Horizons (All measures in percent)

investing for longer horizons generates higher returns on average but also exposes us to more risk. The probability of losing money decreases over time. In fact, there has yet to be a 20-year period in which the S&P 500 index lost money. The last column presents the average return conditional on the return being negative (i.e., the expected loss when a loss occurs).

This last column is an important complement to the probability of a loss, which tells an incomplete story. An investment may have a relatively high probability of loss, but if the average size of our loss is small then this investment may pose less risk than one with a low probability of a loss coupled with a large average loss.

There are several lessons to come from this discussion. First, time diversification is not real. Longer investments, like larger investments, correspond to larger risks. Second, volatility is just one measure of risk and is an imperfect measure. We discussed above that covariance is the relevant measure of risk when considering the addition of an asset to a well-diversified portfolio. Additionally, volatility is a symmetric measure of risk; it captures fluctuations in returns on both the downside and upside. But, large positive returns don't correspond to risk in most circumstances. Investors are primarily worried about *downside* risk - large negative returns. The probability of loss is one way of capturing downside risk, but it too is limited in its usefulness because investors don't just care about the probability of losing money. They also care about how much money will be lost when a loss occurs.

Finally, despite the arguments above, it is still tempting to look at table 5 and think: "If we just invest long enough, we can't lose!" It is true that anyone who has left their money in the stock market for at least 20 years over the last century has never lost money over their investment horizon. However, this only sounds good if our alternative was not investing at all; i.e., our opportunity cost was zero. It is not.

Table 6 presents results from investing in 10-year Treasury notes analogous to those in

Table 5. We would never have lost money investing in a 10-year Treasury assuming we're willing to hold it to maturity. And, our average loss when we do lose on shorter-term investments is relatively small.

Horizon (Years)	Average Return	Return Volatility	Probability of Loss	Average loss when loss
1	5.11	7.64	19.15	-3.86
2	10.59	11.22	7.45	-3.23
5	29.17	24.54	1.06	-1.84
10	70.70	56.63	0.00	N/A
20	216.36	175.68	0.00	N/A

Table 6: 10-Year Treasury Note Performance Statistics Over Different Investment Horizons (All measures in percent)

Perhaps more interesting is that investing in 10-year Treasury notes outperformed the S&P 500 from 1929 to 1948. Admittedly, this was the only 20-year period in which this occurred but broader point remains. Investing over longer time period comes with greater risk.

9.7 Key Ideas

While this chapter has been a little more technical than others, it's important not to lose sight of the key take-aways.

- As investors, we'd like as much reward (expected return) and as little risk (standard deviation) as possible. This conclusion assumes that the only risk we care about is what's measured by standard deviations, and we're **risk averse** - that is we prefer less risk all else equal.
- We can increase expected returns without increasing, and sometimes even with decreasing, the volatility of our investment if we hold multiple assets (i.e., a portfolio of assets). The key is that these assets cannot be too strongly positively correlated. When assets are weakly (or negatively) correlated, combining them diversifies or reduces risk; hence the old adage, "don't put all your eggs in the same basket."
- The optimal portfolio of risky assets has the highest Sharpe ratio, which measures the excess return on an investment (expected return minus risk-free return) per unit of

that investment's risk (standard deviation). Under some assumptions, the portfolio with the highest Sharpe ratio contains a combination of the risk-free asset and the tangency portfolio. *Portfolios of the risk-free asset and the tangency portfolio offer the highest return per unit risk and, as such, are strictly preferable to all other portfolios when investors only care about expected returns and volatility.*

- Which portfolio we want to hold on the efficient frontier is determined by our **risk tolerance**. More risk tolerant investors - people that can stomach more volatility in their savings - can invest less in the risk-free asset and more in the tangency portfolio. Less risk tolerant investors should do the opposite.

These takeaways can also be summarized with two pictures shown in figures 9.17 and 9.18. In both figures, the horizontal axis measures volatility, the vertical axis expected return.




Figure 9.17: Mean Variance and Efficient Frontier of Risky Assets




Figure 9.18: Mean Variance and Efficient Frontier of Risky and Risk-free Assets

In figure 9.17, the peach shaded region contains all of the portfolios consisting of *only* risky assets. The red and blue curve is the frontier of these risky asset portfolios. The red portion of the curve is the efficient frontier of risky portfolios, the blue is the inefficient frontier of risky portfolios. If we could only invest in risky assets - assets whose standard deviations are greater than zero - then we would want to find a portfolio on the red curve. These portfolios offer the highest return for a given level of risk (volatility).

Figure 9.18 introduces a risk-free asset into the investment opportunity set. The risk-free return has no risk - zero volatility - and its return is identified by the purple square on the vertical axis labeled r^F . The portfolios in which we can invest now include the expanded peach shaded region, the blue curve representing the frontier (efficient and inefficient) of risky portfolios, and the red and green lines representing the frontier of all portfolios. The green line is the inefficient frontier, the red line is the efficient frontier.

The efficient frontier contains portfolios consisting of two assets: (1) the risk-free asset (e.g., T-bill) and (2) the tangency portfolio indicated by turquoise circle. I realize the tangency portfolio can contain many assets and so is itself a portfolio. The point is that investors should only invest in some combination of the risk-free asset and this tangency portfolio because there is no way to achieve a higher return without increasing risk. And similarly, there is no way to reduce risk without sacrificing return.

The remaining question at this point is: What is the tangency portfolio in practice when we can invest in many (tens of thousands) assets? In other words, to exploit what we've learned in this chapter, we need to know the tangency portfolio for *all* assets in which we can invest. That's next.

9.8 Technical Appendix

More details on the math and some results from this chapter are presented here.

9.8.1 2-Asset Portfolios

Consider two risky (i.e., volatility > 0) assets, A and B.

Minimum Variance Portfolio of Risky Assets

The variance of a portfolio of assets A and B is

$$Var^p = (w^A)^2 Var^A + (1 - w^A)^2 Var^B + 2w^A(1 - w^A)Cov^{A,B}.$$

To find the portfolio with the smallest variance, and therefore the smallest standard deviation, we take the derivative of the portfolio variance with respect to the portfolio weight w^A , set the derivative equal to zero, and solve for the portfolio weight.

$$\frac{\partial Var^p}{\partial w^A} = 2w^A Var^A - 2(1 - w^A) Var^B + 2Cov^{A,B} - 4w^A Cov^{A,B} = 0$$

Solving for w^A produces equation 9.8. The portfolio weight for asset B is $1 - w^A$.

Tangency Portfolio

The tangency portfolio is the same as the maximal Sharpe ratio portfolio. Consider two risky assets, A and B, that comprise a portfolio, P, and a risk-free asset with return r^F .

The portfolio Sharpe ratio is

$$\frac{r^P - r^F}{SD^P} = \frac{w^A r^A + (1 - w^A) r^B - r^F}{\sqrt{(w^A)^2 (Var^A)^2 + (1 - w^A)^2 (Var^B)^2 + 2w^A(1 - w^A) Cov^{A,B}}}$$

To find the tangency portfolio, we need to maximize this ratio with respect to the portfolio weight, w^A . This can be done by taking the partial derivative of the portfolio Sharpe ratio with respect to the portfolio weight, setting this derivative equal to zero, and then solving for w^A . Computing the derivative is tedious but just repeated application of the chain rule. The result of this process is as follows.

$$\begin{aligned} w^A &= \frac{(r^A - r^F) Var^B - (r^B - r^F) Cov^{A,B}}{(r^A - r^F) Var^B + (r^B - r^F) Var^A - (r^A - r^F + r^B - r^F) Cov^{A,B}} \\ w^B &= 1 - w^A \end{aligned}$$

9.8.2 N-Asset Portfolios

Imagine now we have N assets. We'll use linear algebra to make the notation and results more compact. Otherwise, the formulas just get ridiculously long. However, don't let this scare you like it did me when I first saw it. Matrix notation is just a compact way to express arithmetic in this setting.

Let's define some mathematical objects. The vectors of portfolio weights (w) and expected returns (μ) for assets 1 through N are defined as

$$\begin{aligned} w &= (w^1, \dots, w^N), \text{ and} \\ \mu &= (r^1, \dots, r^N]. \end{aligned}$$

The covariance matrix of returns is

$$\Sigma = \begin{pmatrix} Var^1 & Cov^{1,2} & Cov^{1,3} & \dots & Cov^{1,N} \\ Cov^{2,1} & Var^2 & Cov^{2,3} & \dots & Cov^{2,N} \\ \vdots & \vdots & \vdots & \dots & \vdots \\ Cov^{N,1} & Cov^{N,2} & Cov^{N,3} & \dots & Var^N \end{pmatrix}$$

The variance of each asset's return is on the diagonal. The covariance between each pair of assets is on the off-diagonals. Because $Cov^{i,j} = Cov^{j,i}$, the covariance matrix is symmetric.

The expected return and variance of our N-asset portfolio is just as we saw in equation [9.3](#)

$$w' \mu = r^P = w^1 r^1 + \dots + w^N r^N \tag{9.13}$$

The variance of our N-asset portfolio is similar to that of our 2-asset portfolio (equation 9.4), just with more variance and covariance terms. Recall equation 9.14.

$$\begin{aligned}
 w'\Sigma w = Var^P &= w_1^2\sigma_1^2 + \dots + w_N^2\sigma_N^2 \\
 &+ 2w_1w_2\sigma_{1,2} + \dots + 2w_1w_N\sigma_{1,N} \\
 &+ 2w_2w_3\sigma_{2,3} + \dots + 2w_2w_N\sigma_{2,N} \\
 &+ \dots \\
 &+ 2w_{N-1}w_N\sigma_{N-1,N}
 \end{aligned} \tag{9.14}$$

9.8.3 Limits to Diversification

With an equal-weighted portfolio, w_1, \dots, w_N all equal $1/N$. The portfolio variance - equation 9.14 - can be written as follows.

$$Var^P = \frac{1}{N} \sum_{i=1}^N \frac{1}{N} \sigma_i^2 + \sum_{j=1; j \neq i}^N \sum_{i=1}^N \frac{1}{N^2} \sigma_{i,j}$$

Define the average variance and covariance as follows.

$$\begin{aligned}
 \overline{\sigma^2} &= \sum_{i=1}^N \frac{1}{N} \sigma_i^2 \\
 \overline{cov} &= \frac{1}{N(N-1)} \sum_{j=1; j \neq i}^N \sum_{i=1}^N \sigma_{i,j}
 \end{aligned}$$

We can rewrite the portfolio variance in terms of the average variance and covariance.

$$Var^P = \frac{1}{N} \overline{\sigma^2} + \frac{N-1}{N} \overline{cov}$$

If we take the limit of this expression as N gets large, we get

$$\lim_{N \rightarrow \infty} Var^P = \overline{cov}.$$

The limit to diversification is governed by the average covariance among the assets in the portfolio.

If we assume that all pairs of assets have a common correlation coefficient ρ and volatility σ , then our portfolio variance can be written as

$$Var^P = \frac{1}{N} \sigma^2 + \frac{N-1}{N} \rho \sigma^2$$

When the assets are uncorrelated, risk eventually disappears as we add more and more assets (i.e., as N gets large). When the assets are perfectly positive correlated, adding assets does nothing to portfolio risk, which always equals the common individual asset volatility. In other words, all risk is non-diversifiable.

There are limitations on the average covariance. It turns out that a covariance matrix is **positive semi-definite**, which, loosely speaking, means the average covariance will always be greater than or equal to zero.

Mean-Variance Efficient Risky Portfolios

To find the mean-variance frontier of risky assets, we can either minimize the portfolio variance for a given expected return, or maximize the expected return given a portfolio variance. Let's focus on the first option.

$$\begin{aligned} \min_w \quad & \frac{1}{2}w'\Sigma w \\ \text{s.t.} \quad & \mu'w \geq m \\ & e'w = 1 \end{aligned} \tag{9.15}$$

The $1/2$ is just a scale factor that has no effect on the solution to the problem. The first constraint says that the expected return on the portfolio μw must be at least as large as some constant m . The second constraint says that the sum of the weights must equal one since e is an N -dimensional vector of ones (i.e., $e = (1, \dots, 1)$).

$$e'w = w^1 + w^2 + \dots + w^N$$

We can map out the frontier by varying the expected return target we're trying to achieve, m . Alternatively, we can rely on the **two fund separation theorem**. It turns out that the solution to the minimization problem above can be written like so.

$$w^{Efficient} = (1 - a)w^1 + aw^2$$

where w^1 and w^2 are any two portfolios on the frontier and a is a real number. Now, this may seem circular since we are saying that to find an efficient portfolio we need to know two other efficient portfolios.

What we're really saying is that we only need to solve the minimization problem twice to obtain two portfolios on the frontier. Once we have those two portfolios, we just have to take linear combinations of them to trace out the entire frontier instead of having to re-solve the minimization problem over and over and over...

Minimum Variance Portfolio of Risky Assets

The minimum variance portfolio is the set of portfolio weights, w , that solves the following minimization problem.

$$\begin{aligned} \min_w \quad & w' \Sigma^{-1} w \\ \text{s.t.} \quad & e' w = 1 \end{aligned}$$

The solution to this program is

$$w_{MVP} = (e' \Sigma^{-1} e)^{-1} \Sigma^{-1} e.$$

Plugging this solution into the equations 9.13 and 9.14, we obtain the expected return and variance of the minimum variance portfolio, respectively.

$$r^{MVP} = \frac{\mu' \Sigma^{-1} e}{e' \Sigma^{-1} e} \quad (9.16)$$

$$Var^{MVP} = \sigma_{mvp} = e' \Sigma e \quad (9.17)$$

Tangency Portfolio

The tangency portfolio is the portfolio of risky assets with the maximum Sharpe ratio and therefore solves the following program.

$$\begin{aligned} \max_w \quad & \frac{\mu' w - r^F}{w' \Sigma w} \\ \text{s.t.} \quad & e' w = 1 \end{aligned}$$

The solution (i.e., the portfolio weights for the tangency portfolio) is

$$w^T Tang = \frac{\Sigma^{-1} (\mu - r^F e)}{e' \Sigma^{-1} (\mu - r^F e)},$$

where w^{Tang} is the vector of portfolio weights for the tangency portfolio.

For two-risky assets, call them 1 and 2, the weight on asset 1 in the tangency portfolio is

$$w_1^{Tang} = \frac{(\mu_1 - r_f) \sigma_2^2 - (\mu_2 - r_f) \sigma_1 \sigma_2 \rho_{1,2}}{(\mu_1 - r_f) \sigma_2^2 + (\mu_2 - r_f) \sigma_1^2 - (\mu_1 - r_f + \mu_2 - r_f) \sigma_1 \sigma_2 \rho_{1,2}}$$

The weight on asset 2 is $1 - w_1^{Tang}$.

Two Fund Separation Revisited

Because the minimum variance and the tangency portfolios are on the mean-variance frontier of risky assets, we can use them to map out the entire frontier according to the two fund separation theorem. In other words, every portfolio on the frontier can be identified

$$w^{Efficient} = (1 - a)w^{MVP} + aw^{Tang}$$

For any target return, m , a equals

$$\frac{m(\mu'\Sigma^{-1}e)(e'\Sigma^{-1}e) - (\mu'\Sigma^{-1}e)^2}{(\mu'\Sigma^{-1}\mu)(e'\Sigma^{-1}e) - (\mu'\Sigma^{-1}e)^2} \quad (9.18)$$

9.9 Problems

9.1 (*Conceptual*) Determine which of the following statements are unconditionally true?

- a. Portfolio weights always sum to one.
- b. Portfolio weights are always between zero and one, inclusive.
- c. A portfolio return must be between the minimum and the maximum of its component returns. For example, if a portfolio had ten stocks whose largest return was 12% and smallest return was -8%, then the portfolio return must lie between -8% and 12%, including these extremes.
- d. Equal-weighted portfolios have the same number of shares of each stock in their portfolio.
- e. Value-weighted portfolios will have allocations that are tilted towards (more heavily weighted) larger stocks, i.e., stocks with larger market capitalizations.
- f. If an investor always wants to maintain a value-weighted portfolio, they never have to rebalance their portfolio once set up as such.

9.2 (*Portfolio weights*) Bandit has \$100,000 to invest and is considering four stocks: Revolution Medicines (RVMD), Dollar General Corporation (DG), Hecla Mining Company (HL), and the Allstate Corporation (ALL). The price and number of shares outstanding as of October 13, 2023 for each stock are provided in the following table.

Stock Ticker	Price	Shares
	per Share (\$)	Outstanding (million)
RVMD	29.91	109.45
DG	111.16	219.48
HL	4.14	617.34
ALL	120.32	261.57

Using this information, answer the following questions.

- What is the market capitalization of each firm?
- If Bandit decides to hold an equal-weighted portfolio of the four stocks, what are the portfolio weights? How much money must he invest in each stock? How many shares of each stock must he purchase?
- If Bandit decides to hold a value-weighted portfolio of the four stocks, what are the portfolio weights? How much money must he invest in each stock? How many shares of each stock must he purchase?

9.3 (*Portfolio rebalancing*) The table below presents stock market information for Earthstone Energy Inc. (ESTE) and Natera Inc. (NTRA). Neither stock paid a dividend between September 1 and October 1 of 2023. As of September 1, 2023, Luna owned a value-weighted portfolio of the two stocks worth \$500,000

Date	ESTE		NTRA	
	Price per Share (\$)	Shares Outstanding (million)	Price per Share (\$)	Shares Outstanding (million)
9/1/2023	20.24	106.33	44.25	119.15
10/1/2023	20.40	106.33	39.39	119.15

Using this information, answer the following questions.

- What is the market capitalization of each stock at the beginning of September and October?
- What are the one-month returns to ESTE and NTRA from 9/1/2023 to 10/1/2023?
- Assuming Luna doesn't rebalance her portfolio by buying or selling shares of ESTE or NTRA, what are the portfolio weights, values, and number of shares of ESTE and NTRA in her portfolio on both 9/1/2023 and 10/1/2023?

- d. Now assume that Luna rebalances her portfolio on 10/1/2023 in response to any performance differential between the stocks in her portfolio. How many shares of each stock must Luna buy or sell to maintain her *original* allocation as of 9/1/2023?

9.4 (*Portfolio values and returns*) Number of shares, share price, and dividend data for a portfolio of Cassava Sciences (SAVA) and Dick's Sporting Goods (DKS) are provided in the following table.

Date	SAVA			DKS		
	Shares	Price per Share (\$)	Dividend per Share (\$)	Shares	Price per Share (\$)	Dividend per Share (\$)
2/28/2023	100.00	24.70		50.00	128.63	
3/31/2023	100.00	24.12		50.00	141.89	1.00
4/28/2023	100.00	23.22		50.00	145.01	

Using this information, answer the following questions.


- What are the price returns and total returns for each stock in March and April 2023?
- As of February 2023, what is the value of the portfolio and what are the portfolio weights?
- Assume we reinvest any dividends into purchasing more DKS stock. How many shares of DKS stock can we purchase at the end of March 2023? What are the portfolio values and weights as of March and April 2023? What are portfolio returns during those months?
- Assume we reinvest any dividends into purchasing more SAVA stock. How many shares of SAVA stock can we purchase at the end of March 2023? What are the portfolio values and weights as of March and April 2023? What are portfolio returns during those months?
- Assume we don't reinvest any dividends into the portfolio. What are the portfolio values and weights as of March and April 2023? What are portfolio returns during those months?
- Using your answers to the previous three questions, which dividend reinvestment strategy was optimal (after the fact) assuming the dividend payment sat in cash if not reinvested?

9.5 (*Portfolio weights and returns*) Jenny held a value-weighted portfolio at the start of 2022. She is trying to understand the portfolio allocation and performance of individual assets in the portfolio; however, the information she has been given is incomplete as indicated by the “?”s in the table below. In addition, Jenny has been told that the portfolio earned 11.61% in 2022 and that Match Group Inc. (MTCH) experienced a -9.20% monthly compounded return during 2022.

Company	Ticker	Price per Share (\$)	Shares	Portfolio Weight	Return (%)
			Outstanding (million)		
APA Corp	APA	33.21	307.26	?	76.50
Valero Energy Corp.	VLO	82.97	353.13	41.36	?
Match Group Inc.	MTCH	?	278.09	?	?

Using this information, fill in the missing data from the table. (Hint: Use what you know about the weights and returns to a value-weighted portfolio.)

9.6 (*Portfolio values and returns*) The figure below plots the realized annual returns to portfolios of Granite Construction Inc (GVA) and EMCOR Group Inc. (EME) from October 2022 through September 2023. On the horizontal axis is the portfolio weight on EME.



Using this information, answer the following questions.

- What was the annual return to EME?

- b. What was the annual return to GVA?
- c. Why is the slope of the line positive?
- d. What is the approximate return to an equal-weighted portfolio of EME and GVA?
- e. Describe the portfolio positions (i.e., weights) in EME and GVA for portfolios with returns less than 2.50%.
- f. Describe the portfolio positions (i.e., weights) in EME and GVA for portfolios with returns greater than 30.94%.
- g. If we had instead plotted the portfolio returns against the weight on GVA, how would the figure differ? Specifically, what would be the y-intercept? The slope?

9.7 (*Portfolio expected returns*) Between 2018 and 2022, Xylem Inc (XYL) and Illinois Tool Works Inc. (ITW) have experienced average monthly returns of 1.23% and 0.91%, respectively.

Using only this information, answer the following questions.

- a. What would we estimate as the expected monthly returns to XYL and ITW in January 2023?
- b. From the information given, can we know what the realized returns to XYL and ITW in January 2023 will be?
- c. What is the expected monthly return to an equal-weighted portfolio of XYL and ITW?
- d. What is the expected monthly return to a portfolio consisting of 30% of our wealth invested in XYL and 70% in ITW?
- e. Construct a line plot showing the returns to portfolios consisting of XYL and ITW as a function of the weight on XYL. The horizontal axis of the figure should show the portfolio weight on XYL, the vertical axis should show the portfolio return. Be sure to label both axes, the y-intercept, and the portfolio consisting of only XYL.

9.8 (*Portfolio expected returns and targeting expected returns*) Continuing from the previous problem, assume we own a portfolio worth \$10,000 and consisting of 20% Xylem Inc. (XYL) and 80% Illinois Tool Works Inc. (ITW). We are considering adding a third stock, Ametek Inc (AME), whose average monthly return between 2018 and 2022 was 1.43%.

Using this information, answer the following questions.

- a. How much money is invested in XYL and ITW in our current portfolio?
- b. What is the expected return to our current portfolio?
- c. If we sell \$2,000 worth of ITW to purchase AME, what would the composition of our portfolio be? (I.e., what are the new portfolio weights?) What is the expected return of our portfolio?
- d. Instead of selling any of our existing shares, how much AME stock do we have to purchase and add to our existing portfolio to ensure an expected portfolio return 1.15%?
- e. If we want to achieve a 2% monthly expected return with the three stocks XYL, ITW, and AME, what portfolio should we hold that minimizes the amount of short-selling?

9.9 (*Portfolio summary statistics*) The table below presents monthly returns for Paypal Holdings Inc (PYPL) and the Walt Disney Co. (DIS).

Date	Return (%)	
	PYPL	DIS
1/31/18	15.8924	1.079
2/28/18	-6.9269	-5.0704
3/29/18	-4.4579	-2.6367

Using this information, answer the following questions.

- a. What are the estimated expected returns for PYPL and DIS?
- b. What are the estimated volatilities for PYPL and DIS?
- c. What is the estimated covariance of PYPL and DIS returns?
- d. What is the estimated correlation of PYPL and DIS returns? Would you describe this as high or low? Explain.
- e. Do you think your answers to the previous questions provide accurate estimates of their intended quantities?


9.10 (*Covariance and correlation*) The table below contains annual returns for Agilent Technologies (A) and Qwest Communications (Q).

	Returns (%)	
	Qwest	Agilent
2019	30.53	6.07
2020	(0.49)	40.01
2021	56.39	60.04
2022	(3.00)	(11.95)
2023	(6.87)	(8.60)

Using this information, answer the following questions.

- Create a scatter plot of the two return series and approximate the average return and volatility for the two stocks, as well as the correlation, by eyeballing the figure.
- What are the average returns for Qwest and Agilent?
- What are the volatilities for Qwest and Agilent?
- What are the covariance and correlation of Qwest's and Agilent's returns?

9.11 (*Portfolio Risk and Return*) The scatter plot below presents the daily returns for Freeport Mcmoran (FCX) and John Wiley and Co. (WLY).

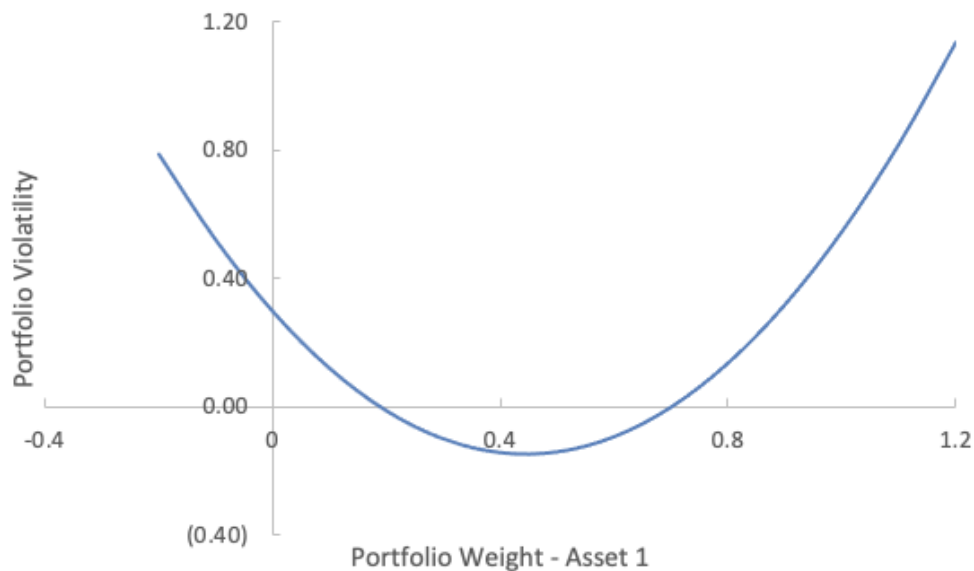


Using this figure, answer the following questions.

- What are the approximate average daily returns for FCX and WLY?

- b. Which stock's returns appear more volatile?
- c. What is the approximate correlation? How would you describe the correlation? Positive or negative? Weak or strong?
- d. What are the approximate coordinates of the red and green points? In other words, what are the FCX and WLY returns corresponding to each of those points?

9.12 (*Portfolio Risk*) Jason, the new analyst at ARM Investors, has constructed the following figure to illustrate the benefits of diversification to a potential client. The chart plots the volatility of a proposed portfolio on the vertical axis against the portfolio weight of one of the component assets on the horizontal axis. What, if any, problems do you see with Jason's figure?



9.13 (*Portfolio Risk and Return*) Lululemon is currently trading at The table below presents information on annual returns for Lululemon Athletic Inc. (LULU) and Apple Inc. (AAPL).

	Returns (%)	
	LULU	AAPL
Average	24.78	29.84
Volatility	37.40	31.39
Correlation	62.87	

Using this information, answer the following questions.

- a. What is the expected return and volatility of a value-weighted portfolio of LULU and AAPL?
- b. What is the expected return and volatility of an equal-weighted portfolio of LULU and AAPL?
- c. How does your volatility answer to the previous question compare to the volatility of either individual stock? Explain the relations.
- d. Would you describe the diversification benefits of portfolios of LULU and AAPL as large or small? Explain.

9.14 (*Diversification*) Determine which of the following statements are unconditionally true?

- a. Diversification means investing in many assets with a low correlation with one another.
- b. Two uncorrelated assets offer greater benefits of diversification than two negatively correlated assets.
- c. We can obtain greater diversification benefits by investing across asset classes (e.g., stocks, bonds, real estate) than within asset classes?
- d. If we can invest in two assets and want to minimize our risk, we should invest all our money in the asset with the lowest volatility.
- e. A well diversified portfolio faces no risk.
- f. Weather events such as hurricanes and earthquakes correspond to systematic risk for insurers.
- g. Defaults in a bank's loan portfolio is an example of idiosyncratic risk.
- h. If asset A offers a higher return and lower volatility than asset B, then we would never want to invest in asset B.
- i. The volatility of a portfolio can never be lower than the lowest volatility of the portfolio's assets if the correlation is positive
- j. Removing assets from a portfolio will typically lead to increased portfolio risk.

9.15 (*Hedge assets*) Tim is a portfolio manager at Citadel whose portfolio has an estimated monthly expected return and volatility of 1.6% and 1.2% per month, respectively. New risk management practices impose a 1% month volatility **risk budget**. To avoid altering the composition of his portfolio, Tim would like to purchase a **hedge asset** -

an asset that is negatively correlated with his portfolio - to reduce his volatility within his risk budget.

Security “Z” has a monthly expected return and volatility of 0.3% and 0.4%, respectively, and a -0.5 correlation with Tim’s portfolio. Assuming Tim has \$750 million invested in his portfolio., answer the following questions using this information.

- How much of the hedge asset must Tim purchase to reduce his portfolio volatility to the target level of 1% per month assuming he receives no new capital, i.e., he is limited to \$750 million and must sell some of his portfolio? What is the expected return to his portfolio inclusive of his hedge?
- How do your answers to the questions in a. change if Tim is able to raise new capital to purchase the hedge asset and leave his \$750 million portfolio unchanged?
- How do your answers to the questions in a. change if the correlation between Tim’s portfolio and security Z is -1.0? 0.5? What do your answers here tell you about the effectiveness of hedge assets?

9.16 (*Portfolio statistics, portfolio allocation*) Bobby Axelrod’s current portfolio allocation is shown in the following table. Expected returns and volatilities are on an annual basis.

	Amount (\$bil)	Expected Return (%)	Return Volatility (%)
Cash	0.25	3.20	0.00
Bonds	2.47	6.89	7.67
Equity	5.89	14.86	18.60

The correlation matrix for Bobby’s portfolio is as follows.

	Cash	Bonds	Equity
Cash	1.00	0.00	0.00
Bonds	0.00	1.00	0.28
Equity	0.00	0.28	1.00

For example, the correlation of cash and bond returns is 0.00, and bond and equity returns is 0.28. The diagonal shows the correlation of each asset’s returns with itself, which is simply 1.00.

Using this information, answer the following questions.

- a. What is the size in dollars of Bobby's portfolio?
- b. What are the portfolio weights?
- c. What is the portfolio expected return?
- d. What is the portfolio volatility?
- e. Bobby is looking to expand into alternative investments with an additional \$1 billion investment. For this purpose, he is considering the following private equity funds whose annual expected returns, volatilities, and correlation with his current portfolio is provided below.

Fund	Expected	Return	Correlation
	Return (%)	Volatility (%)	
AP1	14.50	16.93	0.24
KKR4	11.43	9.36	0.60
SL2	16.71	17.57	0.92
TPG1	10.74	8.64	0.06

In which fund would you recommend he invest if he wants to increase his expected return without increasing his risk? Explain your reasoning without any calculations.

- f. Assuming Bobby follows your answer to the previous question, estimate the portfolio weights, expected return, and volatility of his new portfolio? Is your advice correct? Does Bobby's new portfolio meet his requirements?

9.17 (*Portfolio allocation, risk*) David Weiss is the portfolio manager for a large pension fund that is currently invested in fixed income and public equity securities. The pension is currently worth \$2.04 billion and has an annual expected return and standard deviation of 12.80% and 18.36%, respectively. The pension board of directors has asked Dave to reallocate the pension assets so that 10% of the pension is allocated to private equity. Dave is considering two funds whose expected return, volatility, and correlation with his current portfolio are provided in the following table.

Fund	Expected	Return	Correlation
	Return (%)	Volatility (%)	
TB2	18.29	19.74	0.24
EC9	14.59	13.73	0.60

Using this information, answer the following questions.


- a. How many dollars of the existing portfolio must Dave liquidate to support a 10% private equity target?
- b. What would the expected return and volatility of the new portfolio be if Dave chooses to invest in TB2?
- c. What would the expected return and volatility of the new portfolio be if Dave chooses to invest in EC9?
- d. Based on your answers to the previous two questions, which fund should Dave select? Explain.
- e. After further analysis, Dave argues that he should be allowed to invest more than 10% in the private equity fund selected in the previous question. Is he correct? Explain.

9.18 (*Portfolio allocation, risk*) Andreea has \$2 million in retirement savings invested in a well-diversified portfolio of stocks and bonds. The annual expected return and volatility of this portfolio is 8.7% and 14.3%. She recently received a \$500,000 bonus in the form of company stock, which has an expected return of 12.2% and volatility of 28.6%. Her company stock has a correlation of 0.52 with her current portfolio.

Using this information, answer the following questions.

- a. What fraction of Andreea's wealth is in her company stock? In her portfolio of stocks and bonds?
- b. What is the expected return and volatility of Andreea's assets - her portfolio of stocks and bonds and her company stock? How does it compare to that her portfolio of stocks and bonds alone?
- c. Andreea is concerned about the risk of her savings and is considering selling some of her company stock and investing the proceeds in Treasury securities, which have an expected return of 3.5% and, for the purpose of this problem, an assumed volatility of 0.0%. How much company stock must she sell and invest in Treasury securities to maintain an overall volatility of no more than 14%? What is the corresponding expected return?

9.19 (*Portfolio allocation, mean-variance frontier*) The figure below presents the mean-variance frontier for a portfolio consisting of the S&P 500 index and an index of BAA-rated bonds.



The data labels correspond to portfolio weights on the S&P 500 index.

Using this information, answer the following questions.

- What is the approximate average return and standard deviation of the S&P 500 index?
- What is the approximate average return and standard deviation of the BAA-rated bond index?
- What are the approximate portfolio weights of the minimum variance portfolio?
- Would a risk-averse investor ever want to short the S&P 500 index? How about the BAA-rate bond index?

9.20 (*MVP of stocks and bonds, Tangency Portfolio*) Conventional wisdom suggests that investors should invest 60% of their money in stocks and 40% in bonds - the so-called “60-40” investing rule. The table below provides annual performance statistics for the S&P 500, a portfolio of Baa-rated corporate bonds, and a 3-month T-bill whose volatility has been assumed to be zero for its use as a proxy for the risk-free rate.

	Expected Return (%)	Return Volatility (%)	Covariance
Corp. Bonds	7.19	7.51	
S&P 500	11.82	19.46	0.0058
T-bill	3.33	0.00	

Using this information, perform the following tasks and answer the following questions.

- a. Construct a mean variance frontier for corporate bonds and the S&P 500. Plot the frontier in a figure with portfolio volatility on the horizontal axis and portfolio return on the vertical axis.
- b. Identify on the frontier the following points: 100% invested in the S&P 500, 100% invested in the corporate bond portfolio, and the tangency portfolio.
- c. Construct the capital allocation line and plot this line in the same figure as your mean variance frontier from question a. How is the capital allocation line related to the mean-variance frontier?
- d. What is the Sharpe ratio of the capital allocation line.
- e. A popular investing wisdom is the “60-40 rule” in which investors place 60% of their money in the stock market and 40% in the bond market. Using the mean-variance frontier you constructed, what is the Sharpe ratio of this portfolio? How does it compare to the Sharpe ratio of the tangency portfolio?
- f. What portfolio on the capital allocation line replicates the return of the 60-40 portfolio, and what is its volatility? How does the volatility of the 60-40 portfolio compare with that of the portfolio on the capital allocation line having the same return?
- g. What portfolio on the capital allocation line replicates the volatility of the 60-40 portfolio, and what is its return? How does the return of the 60-40 portfolio compare with that of the portfolio on the capital allocation line having the same volatility?

Part III

Decisions Business People Make

Chapter 10

Estimating the Cost of Capital - Factor Models

Fundamental value relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

We've seen in previous chapters several approaches for estimating expected returns - r in our fundamental value relation.

- the arithmetic average of historical returns,
- the dividend discount and total payout models for stocks, and
- a decision tree for defaultable bonds.

This chapter presents another approach for estimating expected return, **factor models**. We'll focus most of our time on the most popular of these models, the **Capital Asset Pricing Model** or **CAPM**.¹ Despite its shortcomings, the CAPM provides a rich framework from which a number of important implications flow.

- Risk should be measured by how sensitive (correlated) an asset's returns are to the returns of an investor's portfolio, as opposed to the volatility of returns.

¹William Sharpe, as in Sharpe ratio, won the 1990 Nobel Prize for formulating the Capital Asset Pricing Model.

- There are two types of risks found in asset returns: those that can be eliminated by diversification and those that cannot. Only the latter compensate investors with a risk premium or return above and beyond a risk-free investment.
- People should invest their money in a low-cost, well-diversified mutual fund or exchange traded fund (ETF) and treasury securities. How much of each depends on one's risk aversion.
- Investment performance can be measured by excess expected returns, or average returns not accounted for by non-diversifiable risk.

Warning: Like all economic models, the CAPM is wrong in that it vastly simplifies reality and has been rejected by many empirical studies.² Yet, it is useful because it helps us conceptualize from where the cost of capital comes and provides a practical, if imperfect, solution to estimating it.

10.1 Investment Implications

The first big result of the CAPM is that everyone should invest in some combination of the risk-free asset and the market portfolio, where the market portfolio is a value-weighted portfolio of all assets in the economy. In practice, it's impossible to a portfolio of all assets in the economy because many are not tradeable or available for purchase. So, we are often forced to use a proxy for this theoretical portfolio, usually a mutual fund or ETF containing a large number of assets. Some examples as of 2024 include Vanguard's Total Stock Market Index Fund ETF or the Morgan Stanley Capital International (MSCI) World Index ETF (ticker symbol XWD).³

However, to obtain this result, the CAPM relies on some questionable assumptions. Specifically,

1. Investors only care about the mean and variance of their investment returns.
2. Investors share the same beliefs about how asset returns will evolve going forward.

²See the article by Fama, Eugene, F., and Kenneth R. French - 2004. "The Capital Asset Pricing Model: Theory and Evidence." *Journal of Economic Perspectives*, 18 (3): 25-46 - for a summary and critique of empirical evidence relating to the CAPM.

³As broad as the stock holdings in these ETFs are, they do not include other asset classes - bonds, real estate, etc - which are important components of the market portfolio in a CAPM sense.

3. Markets are “perfect” or frictionless.

Assumption 1 is inaccurate but not crazy. Most investors worry about how much they’ll earn on average and how much varies from period to period. However, some (many) investors worry about other features of returns such as the likelihood of a large loss or large gain. Most investors are also more concerned with how much they can lose than how much they can gain and volatility does not distinguish between the two. Assumption 2 is crazy. Most people have different views about how different stocks, bonds, etc. will perform in the future. Assumption 3 is crazy. **Frictionless** or **perfect** markets are hypothetical markets in which there are no taxes or transaction costs, all assets can be traded at any price and in any amount, investors face the same prices for the same assets, and all investors are equally well-informed. Of course, there are taxes and transaction costs when engaging in financial markets. There are also limitations on short-selling assets, and retail investors - most of us - cannot buy and sell assets for the same prices as large institutional investors. Finally, I doubt we’re - readers of this book - are as well informed about most assets as professional investors.

So, taking the CAPM with a grain of salt, the model implies every investor faces the same investment opportunity set (shaded region in figure 10.1), has the same mean-variance frontier of risky assets (the blue curve), and faces the same risk-free rate (the purple square). In other words, everyone is looking at the same picture in figure 10.1. But, if that’s true, then everyone has the same tangency portfolio (the turquoise circle) and the same efficient frontier (the red ray). Everyone wants to hold a portfolio of the risk-free asset and the *same* tangency portfolio. and everyone has access to this capital allocation line.

What differs across investors is the proportion of money allocated to the risk-free asset and the tangency portfolio - where they want to reside on the efficient frontier. This location is dictated by their risk tolerance. More risk averse investors will hold portfolios containing more of the Treasury security and therefore will be on the red line closer to the purple square. More risk tolerant investors will hold portfolios containing more of the tangency portfolio and therefore will be on the red line closer to the turquoise circle.

Regardless of which combination of Treasury security and tangency portfolio investors hold, the fact that everyone is holding the same portfolio of risky assets - the tangency portfolio - implies that this portfolio is the **market portfolio** - a value-weighted portfolio of all assets in the economy. And, because the market portfolio is the tangency portfolio, the market portfolio is mean-variance efficient. Hence, the capital allocation line running through the market portfolio in a CAPM setting is referred to as the **capital market line** or **CML**.




Figure 10.1: Mean Variance and Efficient Frontier of Risky and Risk-free Assets

This logic is how we arrive at the CAPM's investment recommendation. We should put all of our savings into a portfolio consisting of a risk-free asset, like a Treasury security, and the market portfolio, such as Vanguard's Total Stock Market Portfolio or MSCI's World Index. No more sorting through individual stocks and bonds to find the golden ticket or the perfect combination of assets. Similarly, actively managed investment funds charging high fees make no sense in a CAPM world. The only thing that might change is how much we invest in the risk-free asset and the market to reflect changes in our risk tolerance. As we age, we may want to shift our holdings from the market to the risk-free asset, assuming our risk tolerance declines with age.

10.2 Estimating Expected Returns

Previously we estimated expected returns, $\mathbb{E}(r)$, with a simple average,

$$\bar{r} = \frac{1}{T} (r_1 + \dots + r_T)$$

The r_1 through r_T correspond to a sample of historical realized returns. The sample average, \bar{r} , is our numerical estimate of expected return - r in our fundamental value relation. Equivalently, it is an educated guess of what an asset, like a bond or stock, might earn in the future.

We also estimated expected stock returns using the Gordon growth and total payout models. Recall that for the former, the expected stock return equals the sum of the dividend yield and dividend growth rate, g .

$$r = \frac{Div_1}{Price_0} + g$$

For the total payout model, expected returns equals the sum of the payout yield and earnings growth, g .

$$r = \frac{d \times Earn_1}{Price_0} + g$$

The fraction of earnings distributed via dividends or share repurchases is denoted by d .

The CAPM provides another model-based approach for estimating expected returns.

$$r = r^F + \beta (r^M - r^F) \quad (10.1)$$

Equation 10.1 says that the expected return on an asset - this could be *any* asset - equals the sum of two numbers.

1. Risk-free return, r^F . This term represents the minimum amount of compensation an investor should earn. In other words, it's compensation for the time value of money when there is no risk in the investment.
2. **Risk premium**, $\beta (r^M - r^F)$. This term represents the additional compensation an investor will earn when taking risk; hence, the name risk premium.

Theoretically, the risk-free return, r^F , is a constant - a number that never changes. In practice, we'll approximate this with the current yield on a Treasury or other near-default-free bond which is hardly risk-free because of changing interest rates (recall chapter 7).

The risk premium is the product of two terms.

1. **Beta**, β . This term represents the asset's **market risk exposure**, or how sensitive the asset's returns are to variation in market returns. Beta answers the question: By how much does an asset's return change when the expected return on the market, r^M , changes? Assets sensitive to changes in expected market returns are relatively risky, have high betas, and offer higher returns. Assets insensitive to changes in expected market returns are relatively safe, have low betas, and offer lower returns.

2. **Market risk premium**, $(r^M - r^F)$. This term represents the compensation investors receive for investing in the market. It is referred to as the **market risk factor** and **excess market return** because it measures the return to the market in excess of a risk-free return.

The product of beta and the market risk premium tells us how much additional return investors should expect for investing in an asset with a certain exposure to market risk. The larger an asset's beta, the more sensitive its returns are to changes in market expected returns, the riskier the asset. Asset's with large betas are said to have large *risk exposures*.

So, beta is a new measure of risk, distinct from standard deviation (a.k.a., volatility). Whereas the standard deviation measures how an asset's return's fluctuate around its average, beta measures how sensitive an asset's returns are to returns to a market portfolio. We'll define this measure more precisely below but for now its useful to have some intuition under our belts.

The question now is: How can we use equation 10.1 in practice? In other words, how do we put numbers to r^F , β , and r^M to get a number for r , the expected return in which we're interested in estimating? As always, an example will be useful.

10.3 Estimating Apple's Equity Cost of Capital

Let's estimate Apple's equity cost of capital - what they're shareholders expect to earn on an annual basis and the discount rate for the cash flows they receive. As a benchmark, let's first estimate this quantity with Apple's average return from 2017 to 2021, which was 3.52% on a monthly basis. Annualized, this figure implies an expected return of 42.28%.

Let's compare this estimate with one we get from using the CAPM. We'll proceed in three steps by estimating each of the components of equation 10.1: r^F , $(r^M - r^F)$, and β .

10.3.1 Step 1: Estimate the Risk-Free Rate

As we've said, there's no such thing as a risk-free asset in the real world so we'll have to use a proxy to estimate r^F . The most common proxy is a U.S. Treasury security. If we were in another country, we might use the yield on that country's sovereign debt (i.e., government bonds). For countries with risky government debt, we might instead use the yield on relatively safe corporate debt. Regardless, we have to select a maturity. For Treasury's,

do we use a T-bill? T-note? T-bond? And, exactly which maturity? 30-day? 5-year? 30-year?

The choice of specific Treasury depends on the horizon of the cash flows being discounted. For example, if we had cash flows stretching out over a 3-month horizon, we would use the yield on a 90-day T-bill. If the cash flows stretched out over a 5-year horizon, we would use the yield on a 5-year T-note. If we had an indefinitely long horizon - say we were trying to value Apple's stock - then we would use the yield on the 30-year T-bond, the longest maturity Treasury.

The yield curve as of December 31, 2021 is presented in Figure 10.2. Let's assume we are valuing Apple's stock and as such need to discount cash flows into the indefinite future (companies don't have an expiration date). In this case, we'll use the yield on the 30-year T-bond, 1.90%, as our proxy for the risk-free rate, r^F .




Figure 10.2: March 19, 2021 Treasury Yield Curve (Source: U.S. Treasury)

10.3.2 Step 2: Estimate the Market Risk Premium

The market risk premium, $(r^M - r^F)$, requires an estimate of the expected return on the market and the risk-free rate. We already have an estimate of the risk-free rate from the current yield on a 30-year T-bond, 1.90%. In practice, we often use a different estimate of the risk-free rate to compute the market risk-premium. In theory, this makes no sense. There is only one risk-free rate and it never changes. In practice, there is a rationale.

Treasury returns vary over time. Our best guess of something that varies is its expectation, which we estimate with an average. So, just like we can take an average of historical market returns to estimate the expected return on the market, we can take an average of historical bond returns to estimate the expected return on the Treasury. The difference between these two estimates gives an estimate of the market risk premium.

Investment	Returns (%)	
	1927-2021	1970-2021
Treasury bills (30-day)	3.30	4.46
Treasury notes (10-year)	5.11	7.24
S&P500	11.82	12.33

Table 1: Average Annual Investment Returns

Table 1 presents historical average annual returns for the U.S. stock market, as represented by the S&P 500 index, the 30-day T-bill, and the 10-year T-note.⁴ Estimates of the market risk premium using data from 1927 to 2021 range from 8.52% to 6.71%, depending on whether we use the T-bill or T-note as a proxy for the risk-free rate. Using data from 1970 to 2021, those estimates vary from 7.87% to 5.09%. Statistically speaking, the range of reasonable estimates is quite large - in practice 3% to 8%. For the purposes of estimating Apple's equity cost of capital, we'll use a 6% of market risk *premium*.

An alternative approach to estimating the market risk premium starts by estimating the expected return on the market using the Gordon Growth Model from chapter 8.

$$Value_0 = \frac{Div_1}{r - g}$$

This model assumes the stock pays a stream of dividends growing at a constant rate, g , forever. Solving for the expected return, r , yields

$$r = \frac{Div_1}{Value_0} + g$$

In words, the expected return on the asset equals the sum of its dividend yield and dividend growth.

We can apply this result to any asset, including a portfolio of assets such as the market portfolio. As of December 2022, the dividend yield on the S&P 500 was 1.74%. If dividends were expected to grow at 6% per year, the implied expected market return would be 7.74%. Depending on the risk-free rate, the corresponding market risk premium, $r^M - r^F$, would be something less.

⁴The U.S. stock market is a value weighted portfolio of all stocks on the NYSE, AMEX, and NASDAQ.

10.3.3 Step 3: Estimate Beta (β)

To estimate beta, we first need to know what it is. We should also be more precise. The beta in equation 10.1 is called a **market beta** because it measures the sensitivity of asset returns to variation in the market return. There are other betas measuring the sensitivity of asset returns to other risk factors. We'll explore these other factors towards the end of this chapter. Unless explicitly stated, beta and market beta will be used interchangeably.

The beta of an asset - call it asset "A" - with respect to the market is defined as follows.

$$\beta^A = \frac{Cov^{A,M}}{Var^M} \quad (10.2)$$

The covariance term in the numerator captures how asset A's returns move in relation to the market's returns. The variance in the denominator is there for scale. Covariance and variance units are both squared returns. So, dividing the covariance by the variance removes these squared return units.

Beta is easily estimated in a spreadsheet or other software program by computing the covariance between asset A and the market returns and the variance of the market return. To do so, we'll need some data. A common choice is the most recent five years of monthly data.⁵ From January 2017 to December 2021, Apple's return covariance with a broad-based U.S. stock market index was 0.0023, and the market index's variance was 0.00206. Apple's beta is therefore

$$\beta = \frac{0.00232}{0.00206} = 1.13.$$

Below and in the technical appendix, we'll look at alternative ways to estimate beta.

What does beta mean? Look at equation 10.1. For a one percent increase in the market risk premium, an asset's cost of capital increases by β percent. For Apple, a one percent increase in the market risk premium is associated with a 1.13% increase in Apple's expected return.

10.3.4 Step 4: Putting it all together

With our estimates for each component of equation 10.1, we can compute the cost of capital for Apple as of December 31, 2021.

$$r^{Apple} = r^F + \beta^i (r^M - r^F) = 0.019 + 1.13 \times 0.06 = 0.0868$$

⁵Another common choice is the most recent two years of weekly data.

Apple's shareholders should expect an 8.68% return on their stock investment each year. Equivalently, it costs Apple 8.68% to raise money from shareholders.

Based on Apple's historical performance, 8.68% seems awfully low especially when compared to his recent historical performance - 3.52% *monthly* average return. How can the value from the CAPM be so far off from Apple's recent realized returns?

First, the CAPM estimates the *expected* return. We know from previous discussions that *realized* returns are often very different from *expected* returns. Second, the realized returns are historical and not necessarily indicative of future returns even though they are helping us predict future returns by providing an estimate for beta and the market risk premium. Third, maybe the model is wrong; maybe there are other risk factors, besides the market risk factor, that matter for expected returns. All of these are plausible explanations, the last of which we'll explore a little more deeply towards the end of this chapter.

10.4 Context for Beta

Is a beta of 1.13 small? Large? It would be nice to have some context to interpret this number. Related, how does the market risk exposure that beta measures relate to the return risk that volatility measures?

10.4.1 Beta for the Market Portfolio = 1

The beta of the market portfolios is one. Mathematically, this fact can be seen in equation 10.1.

$$r = r^F + 1 \times (r^M - r^F) = r^M$$

It can also be seen in equation 10.2. The beta for the market portfolio, β^M is

$$\beta^M = \frac{Cov^{M,M}}{Var^M} = \frac{Var^M}{Var^M} = 1. \quad (10.3)$$

In the numerator, we recognized that the covariance of the market return with itself equals the variance of the market return. This fact holds more broadly. The covariance of any random variable with itself is the variance of that random variable. Thus, investing in the market portfolio - a broadly diversified portfolio of assets - exposes investors to a beta equal to one.

Betas greater than one correspond to assets whose returns are relatively sensitive to market fluctuations and, as such, offer higher returns than the market, on average. Betas less

than one correspond to assets whose returns are relatively *insensitive* to market fluctuations and, as such, offer lower returns than the market, on average.

Now, we need to be careful. Low beta does *not* mean no variation! An asset can have a returns that are largely insensitive to what the broader economy is doing but still have returns that are quite volatility. In other words, the asset's returns will bounce around over time, reflecting volatility. The low beta just tells us that the bouncing around is largely unrelated to when the market's returns are bouncing around. Figure 10.3 illustrates just such a situation.




Figure 10.3: High Volatility, Low Beta Returns

The figure presents monthly returns from January 2017 to December 2021 for a value-weighted portfolio of several thousand U.S. stocks and Concord Medical Services (CCM), a healthcare company. The figure shows that, relative to the market, Concord's stock returns are quite volatile as suggested by the different scales for the two series. In fact, the monthly return volatility of Concord stock is 13.8%, more than three times as volatile as the stock market at 4.5%.

Despite its volatility, Concord's beta is 0.14 suggesting that Concord's stock returns are only weakly related to the market returns. Indeed, the ups and downs in Concord's stock returns don't seem to coincide with the ups and downs in the market. Why? Because whether the economy is doing well or poorly, consumers need the medical services Concord's companies provide.

Contrast this situation with that presented in figure 10.4, which replaces Concord's return series with that of Advanced Micro Devices (AMD), an integrated chip maker. AMD's volatility is similar to Concord's, 16.1% per month, significantly higher than that of the market. However, AMD's beta is 1.76 indicating a great deal of sensitivity to the market. When the market moves up, AMD tends to move up by a lot and similarly when the market moves down. Indeed, the wiggles of the market and AMD shown in the figure tend to overlap. Why? Because demand for AMD's computer chips are particularly sensitive to business conditions, much more so than the demand for healthcare services.




Figure 10.4: High Volatility, High Beta Returns

10.4.2 Beta for the Risk-free Asset = 0

Any asset with a risk-free return has a beta equal to zero. There are two ways to see this. First, plug zero into equation 10.1.

$$r = r^F + 0 \times (r^M - r^F) = r^F$$

Now go back to the definition of beta in equation 10.2. The numerator is the covariance of the asset's return with the market return. But, a risk-free asset's return doesn't vary (in theory). So, the covariance of something that doesn't vary (risk-free return) with something that does (the market return) is zero. Hence, the risk-free asset has a beta of zero and its expected return according to the CAPM is the risk-free rate.

Notice, that any asset could have an expected return equal to the risk-free rate as long as its returns were uncorrelated with those of the market.⁶ We could have an asset whose returns are highly volatile (big standard deviation) but that are uncorrelated with the market (i.e., beta equals zero). Therefore, this asset should offer a risk-free return, despite its large volatility. Such assets are referred to as **zero-beta assets**. This may sound odd - an asset with a lot of volatility only offering a return equal to the risk-free rate. Who would hold such an asset and why? The answer lies in distinguishing between different types of risk.

10.4.3 Systematic Risk versus Idiosyncratic Risk

The CAPM distinguishes between two types of risk.

1. **Market.** Also known as **non-diversifiable**, **systematic**, and **priced** risk, this risk affects all assets and cannot be reduced by diversifying our investments. Examples of systematic risks affecting all assets include changes in central bank policies that alter interest rates, changes in fiscal (e.g., tax) policies, wars, and recessions.
2. **Idiosyncratic.** Also known as **diversifiable** or **firm-specific** risk, this risk is specific to an asset or collection of assets and can be eliminated by diversifying our investments. Examples of idiosyncratic risk include the death of a CEO, a warehouse fire, a labor strike, and a product failure. Intuitively, when the price of one stock goes down because of an idiosyncratic event, the effect is small and often offset by price increases of other stocks in a diversified portfolio.

Systematic risk is captured by the risk premium term in equation 10.1, $\beta^i (r^M - r^F)$. Idiosyncratic risk is a little more subtle.

Let's write an asset's return - not expected but actual - as follows.

$$r = r^f + \beta^i (r^M - r^F) + e \quad (10.4)$$

This equation says that the return on an asset is equal to the risk-free rate plus the risk premium plus an error, e , that has zero mean and is uncorrelated with the market and risk-free returns. The terms r , r^M , and e are all random variables that can take on many different values with different probabilities. This equation will always hold no matter what the risk-free rate and return to the market are because the error term ensures it's true.

⁶Remember from chapter 9 that two returns that are uncorrelated (i.e., zero correlation) have zero covariance.

Taking expectations of both sides of equation 10.4 recovers equation 10.1 because our best guess for the error term e is its expectation, zero.

If we take the variance of both sides of equation 10.4 we can see from where an asset's risk comes.

$$Var(r) = \underbrace{\beta^2 Var(r^M)}_{\text{Market risk}} + \underbrace{Var(e)}_{\text{Idiosyncratic risk}} \quad (10.5)$$

Equation 10.5 shows that return variation - why asset returns bounce around over time - comes from two sources. The first is from the asset's exposure to market risk. If the firm has no exposure to market risk - beta equals zero - then it faces no market risk. The second source is from random events that are uncorrelated with the market, e .

So, beta measures a very specific component of risk, namely, systematic risk. Our volatility measure - which is the square root of the left side of equation 10.5 - captures *both* systematic and idiosyncratic risk. This is the key difference between these two measures of risk and it's worth repeating. Beta measures systematic or market risk; volatility measures total risk which equals market risk plus idiosyncratic risk.

Consider Apple. The variance of Apple's monthly returns from January 2017 to December 2021 is 0.0072. The variance of the market return over this period is 0.0021. Rearranging equation 10.5 gives us the error variance.

$$Var(e) = Var(r) - \beta^2 Var(r^M) = 0.0072 - (1.13)^2 \times 0.0021 = 0.0046$$

In other words, most of Apple's stock return variation comes from idiosyncratic risk - 63.7% (0.0046/0.0072). The remaining 36.3% (0.0026/0.0072) comes from its exposure to market risk. These proportions are not uncommon among individual stocks whose returns vary primarily because of idiosyncratic events, as opposed to exposure to market risk.

Notice that the e in equation 10.4 does not appear in the equation for expected returns 10.1. As we said above, on average this error equals zero so its expectation is equal to zero. In other words, *idiosyncratic risk does not affect expected returns according to the CAPM*. This is important. Random events uncorrelated with the market affect stock prices, but do *not* affect what we should expect to earn on our investments. This seems almost contradictory! The key to understanding this fact is that investors can costlessly eliminate idiosyncratic risk by holding diversified portfolios according to the CAPM. Idiosyncratic risk across a large number of assets cancels each other out. Some stocks have bad news, others good. As such, idiosyncratic risk according to the CAPM is "not priced" and is referred to as **unpriced risk**.

In contrast, the market risk premium *does* appear in equation 10.1. Therefore, the risk that comes from variation in the market return - the systematic risk - is “priced” or referred to as **priced risk**. It is risk that affects expected returns because no matter how well-diversified investors portfolios are, we cannot get rid of the risk that affects *all* stocks. How expensive that market risk is for an asset depends on the asset’s risk exposure or beta.

10.4.4 Beta Variation

How much do betas vary over time and across assets? Figure 10.5 shows Apple’s and AMD’s estimated equity beta’s between December 2004 and December 2021. Each estimate is computed using the previous five years of monthly data. These estimates are referred to as **rolling betas** because they are constructed using a rolling window of data.⁷




Figure 10.5: Time Variation in Market Betas

The figure shows that betas can vary a lot over time. Apple’s beta varies from a high of 2.50 in September 2008 to a low of 0.75 in April 2015. For comparison, AMD’s beta varies from a high of 3.85 in June 2006 to a low of 1.75 in December 2021. This difference, or **range** in statistical parlance, of 2.1 translates into a big difference in the equity cost of capital for AMD. (Remember that expected returns and cost of capitals are the same thing.)

⁷In other words, we start by estimating beta using data from January 2005 to December 2009. Then we use data from February 2005 to January 2010 to estimate a new beta. And so on until we get to the last window of data from January 2017 to December 2021.

Assume interest rates and market risk premium are constant at 1.9% and 6%, respectively. At an equity beta of 1.75, AMD's equity cost of capital is $1.9\% + 1.75 \times 6\% = 12.4\%$ according to equation 10.1. At an equity beta of 3.85, AMD's equity cost of capital is $1.9\% + 3.85 \times 6\% = 25\%$. That's a difference of 12.6%! That's a big number whether it's from the perspective of the investor and what they should expect to earn in a year (and the risk they're taking), or from the perspective of AMD's management that needs to raise money to fund its investments.

When determining which beta to use, it might seem obvious to use the most recent. However, the relevant beta is the one that most accurately captures the risk exposure *in the future*. If a company has a large restructuring or investment planned, the most recent beta may be inappropriate. In these cases, it may make sense to use the beta from other, comparable companies. However, there are some subtleties with doing so that are postponed until chapter 12.

Figure 10.6 zooms out a bit by presenting average equity betas across industries as of December 2022.⁸ The industries are sorted by their beta so that low beta industries are on the left side of the figure, high beta industries on the right.

The lowest beta industries are tobacco products and utilities (electric, gas, and water supply). This makes sense. Tobacco is an addictive product and good times or bad, smokers are going to smoke. Likewise, people need basic utilities regardless of the current state of the economy. In other words, just because the economy is doing poorly doesn't mean we shut off all the electric and gas. In the Northeast U.S., you have to heat home in Winter, and in the Southwest U.S., you have to cool your home in the summer. That isn't to say that these industries are completely immune to economic conditions. Their betas are not zero. In bad times perhaps smokers smoke fewer cigarettes, and we turn down the heater/air conditioning a little in the Winter/Summer.

In contrast, oil and gas extraction and lumber and wood products are the highest beta industries. This too makes sense. The profitability of oil and gas extraction is largely dependent on oil and gas prices, which in turn are a function of the state of the economy. When the economy is doing well, consumers travel more and demand more fuel. Lumber and wood are the primary input to construction, which is another industry that is very sensitive to how well the economy is doing. However, unlike construction services which can potentially shift from new construction to maintenance and repair, raw materials have less flexibility. Both industries are also highly capital intensive, facing large fixed costs that increase their sensitivity to economic swings.

⁸The betas were estimated using the previous five years of monthly data.




Figure 10.6: Market Betas by Industry

10.4.5 Debt Betas

The CAPM applies to *all* assets, not just stocks, though most applications of the model are on equity securities. For example, we can estimate the debt cost of capital using the CAPM. Recall from chapter 7 that the yield on a bond is equal to the expected return of the bond only if the bond is free from default risk. Of course, most (all) bonds face some degree of default risk and therefore the expected return on most bonds requires estimation. While the Technical Appendix to chapter 7 provides one estimation approach, the CAPM provides another.

The challenge with applying the CAPM to bonds, or other illiquid assets, is that they don't trade very often and therefore historical returns can be unreliable. That is, it's difficult to estimate the bond beta we need to plug in to our CAPM result, equation 10.1. We can however, use bond index data to estimate betas for groups of bonds distinguished by credit rating and maturity. Table ?? presents debt betas by credit rating and maturity averaged

across industries. The betas by maturity are based on bonds rated BBB and above (i.e., high credit quality, low default risk).

Credit Rating	Beta	BBB and above	
		Maturity (Years)	Beta
A and above	< 0.05	1-5 year	0.01
BBB	0.10	5-10	0.06
BB	0.17	10-15	0.07
B	0.26	> 15	0.14
CCC	0.31		

Table 2: Debt Betas by Credit Rating and Maturity. Source: Shaefer and Strebulaev, 2009, “Risk in Capital Structure Arbitrage,” Working Paper, Stanford GSB

The results in table ?? are consistent with intuition. Lower rated and longer maturing bonds tend to be riskier as suggested by their higher betas. The higher betas, in turn, imply higher expected returns. Yet, even the riskiest bonds rated “CCC,” have a beta that is less than the typical equity beta. (See figure 10.6).

To clarify, the debt betas do *not* imply that borrowers’ requirements to repay their debts depend on the state of the economy. Borrowers are always required to repay their debts regardless of how well the economy is doing. The positive betas imply that either the default risk, the recovery rate, or both tend to vary with the state of the economy. That is, firms are more likely to default in bad states of the economy, or the amount creditors recover in default tends to be lower in bad states of the economy.

10.4.6 Portfolio Betas

Like duration, the beta of a portfolio is weighted average of the component betas. Mathematically,

$$\beta^P = w^1\beta^1 + w^2\beta^2 + \dots + w^N\beta^N \quad (10.6)$$

where the weights are value-weights. Let’s look at an example.

Table 3 presents information on a portfolio of three stocks. The weights are computed by dividing the amount invested in the stock by the total amount invested in the portfolio. For example, we invested \$300 in Dish Network implying that its portfolio weight is $300/1,000 = 0.30$. The portfolio beta is computed using equation 10.6.

$$\beta^P = 0.20 \times 4.95 + 0.30 \times 1.66 + 0.50 \times 2.66 = 2.82$$

Stock	Position (\$)	Weight	Beta
QuantumScape (QS)	200	0.20	4.95
Dish Network (DISH)	300	0.30	1.66
Palantir Technologies (PLTR)	500	0.50	2.66
Portfolio	1,000	1.00	2.82

Table 3: Portfolio Beta Calculation

If we think about it, this is a really useful result. It means we can achieve almost any market risk-exposure (i.e., beta) we want through a judicious choice of assets and positions (i.e., long or short). Consider the portfolio detailed in table 4. With the same amount of money, \$1,000, we can achieve a 0-beta or **market neutral** portfolio, a portfolio with no exposure to movements in the market. And, according to the CAPM, the expected return of this portfolio should equal the risk-free rate.

Stock	Position (\$)	Weight	Beta
QuantumScape (QS)	-681	-0.68	4.95
Dish Network (DISH)	1,100	1.10	1.66
Palantir Technologies (PLTR)	581	0.58	2.66
Portfolio	1,000	1.00	0.00

Table 4: Portfolio Beta Calculation with Short Position

10.5 Bringing it all Together

Figure 10.7 brings most everything we've learned in the previous and current chapters together. Let's walk through it slowly to highlight and connect concepts.

10.5.1 Capital Market Line and Efficient Portfolios

The left plot in figure 10.7 presents the mean-variance frontier of 12 risky assets (black curve) and the efficient frontier (red ray) that we analyzed in chapter 9.⁹ For the purpose of our

⁹The 12 risk assets are: Microsoft (MSFT), Archer Daniels Midland (ADM), International Business Machines (IBM), Hershey (HSY), General Mills (GIS), Proctor & Gamble (PG), Caterpillar (CAT), Deere & Co (DE), Boeing (BA), JP Morgan Chase & Co (JPM), Wal-Mart Stores Inc (WMT), and EBAY Inc (EBAY).




Figure 10.7: Capital Market Line (CML) and Security Market Line (SML)

discussion, we'll treat “the market” as consisting of these twelve assets. Expanding the set of assets to include all stocks, bonds, real estate, etc. has no impact on the lessons.

The horizontal axis in the left figure measures standard deviation; the vertical axis measures expected returns. We estimated both statistics using the standard deviation and average of historical returns. The figure also highlights several specific portfolios, where the term portfolio is used in a very general sense to also include holdings of a single asset - i.e., portfolios placing 100% weight on one asset and 0% weights on all other assets. The different portfolios are labeled in the figure, but it's worth pointing them out to be perfectly clear.

- Risk-free asset - purple square on the vertical axis.
- Market portfolio - turquoise circle. (This is also our tangency portfolio according to the CAPM.)
- Minimum variance portfolio of risky assets (MVP) - red diamond.
- Boeing Company (BA) - blue circle in the interior of the risk-asset frontier.
- Efficient portfolio 1 - green circle on the red ray that is mean-variance efficient because it lies on the efficient frontier.
- Efficient portfolio 2 - orange circle on the red ray that is mean-variance efficient because it lies on the efficient frontier.

There are several new insights in this figure.

Finding Efficient Portfolios

Mean-variance analysis tells us that the best we can do - highest return for a given level of risk, or lowest risk for a given level of return - is invest in portfolios comprised of the risk-free asset and the tangency portfolio, i.e., portfolios on the efficient frontier. The CAPM tells us that the tangency portfolio is the market portfolio. So, for every asset in the economy, there is a corresponding efficient portfolio in which we would be better off investing. Let's use our 12-stock frontier and Boeing stock as an example. In other words, the "market" for this example consists of the 12 stocks, one of which is Boeing.

Efficient portfolio 2 (orange circle) answers the question: What portfolio offers the same risk as Boeing (BA) but with the most reward? Starting from point BA, if we draw a vertical line, thereby maintaining the same standard deviation, Efficient Portfolio 2 is the portfolio offering the highest expected return.

What is this portfolio? We know it has the same standard deviation as Boeing, 19.9%. In other words,

$$SD^{BA} = \sqrt{(w^F)^2 Var^F + (w^M)^2 Var^M + 2w^F w^M Cov^{M,F}} \quad (10.7)$$

where superscripts for F and M correspond to the risk-free asset and market portfolio. Because the variance of the risk-free asset equals zero, the first term in equation 10.7 is equal to zero. Similarly, because the covariance of the risk-free asset with any asset is equal to zero, the third term in equation 10.7 is equal to zero. What's left is the following.

$$SD^{BA} = \sqrt{(w^M)^2 Var^M} = w^M SD^M \quad (10.8)$$

More generally, the standard deviation of *any* portfolio on the efficient frontier equals the weight on the market, w^M , times the standard deviation of the market, SD^M . Therefore, the weight on the market in the efficient portfolio is equal to the ratio of volatilities.

$$w^M = \frac{SD^{BA}}{SD^M}$$

(We just divided both sides of equation 10.8 by SD^M .) Plugging numbers estimated from historical returns produces a portfolio weight on the market portfolio equal to

$$SD^{BA} = 0.199 = w^M SD^M = w^M 0.162 \implies w^M = \frac{0.199}{0.162} = 1.22.$$

What did we just do and what does it mean? Boeing stock offers us a risk-reward profile of 19.9% volatility for a 20.6% expected return. If we instead invest 122% of our wealth in

the market portfolio and -22% in the risk-free asset (i.e., short a Treasury, borrow money at the risk-free rate), we would face the exact same volatility as investing in Boeing stock - 19.9%. However, our expected return would be.

$$w^F r^F + w^M \mathbb{E}r^M = -0.22 \times 0.031 + 1.22 \times 0.3325 = 0.3988,$$

or 39.88%, significantly higher than the 20.6% offered by Boeing. Thus, investing in efficient portfolio 2 (the orange circle) offers a much more appealing risk-reward profile than investing in Boeing stock alone.

Efficient portfolio 1 (green circle) answers the question: What portfolio offers the same expected return as Boeing (BA) but with the least risk? Starting at point BA, if we draw a horizontal line, thereby maintaining the same expected return, Efficient Portfolio 1 is the portfolio offering the lowest risk.

What is this portfolio? We know it has the same expected return as Boeing, 20.6%. So, we can set the expected return to the efficient portfolio - risk free asset and market portfolio - equal to Boeing's expected return and solve for the market portfolio weight, w^M , like so.

$$\mathbb{E}(r^{BA}) = 0.206 = (1 - w^M)r^F + w^M \mathbb{E}(r^M) = (1 - w^M)0.031 + w^M 0.3325 \implies w^M = 0.582$$

If we want to earn 20.6% per year, we would be much better off investing 58.2% of our wealth in the market and 41.8% of our wealth in the risk-free asset. The standard deviation of efficient portfolio 1 is

$$w^M SD^M = 0.582 \times 0.1622 = 0.094,$$

or 9.4%, significantly lower than the 19.9% that Boeing stock experiences.

What this discussion highlights is that it doesn't make sense - according to the CAPM - to invest in anything other than the market portfolio and a risk-free asset. For any stock or group of stocks *not* on the efficient frontier, we can always find an efficient portfolio with the same volatility and higher expected return and another efficient portfolio with the same expected return and a lower volatility.

Risk Decomposition

The risk decomposition in equation 10.5 is shown visually along the horizontal axis, again using Boeing as an example. The total risk of Boeing's returns is identified by its volatility,

which is 19.9% per annum. We know from equation 10.5 that this risk can be decomposed into two pieces: systematic and idiosyncratic.

The idiosyncratic piece is the risk that can be eliminated by diversification - by holding the most efficient portfolio possible. For Boeing's 20.6% expected return, the most efficient portfolio is Efficient Portfolio 1 (green dot), which has the same return but significantly less risk. The horizontal distance between Efficient Portfolio 1 and Boeing (point BA) corresponds to idiosyncratic or diversifiable risk. The horizontal distance between Efficient Portfolio 1 and the vertical axis corresponds to systematic or non-diversifiable risk. For an expected return of 20.6%, we cannot reduce risk any further than that of Efficient Portfolio 1.

Numerically, the risk for Boeing stock breaks down as follows.

- Total risk = 20.6%. This is just the standard deviation of historical returns.
- Systematic risk = 9.4% (We computed the volatility of Efficient Portfolio 1 just above.) This is the volatility of $\beta^{BA}(r^M - r^F)$, where β^{BA} is Boeing's equity beta.
- Idiosyncratic risk = 11.2% (20.6% - 9.4%). This is the residual risk or risk of the error term, e , in equation 10.4.

Investing in Boeing *alone* is inefficient in this setting because we can get the same expected return with much less risk by diversifying our investments and investing the market portfolio and the risk-free asset (e.g., Treasury security).

Capital Market Line

If we accept the assumptions of the CAPM so that the tangency portfolio is the market portfolio, then the efficient frontier is called the **capital market line** or **CML**. Remember, this line contains only efficient portfolios - portfolios consisting of the risk-free asset and the market portfolio - nothing else. The equation of this line is just a variation of what we saw in chapter 9.

$$r^P = r^F + \underbrace{\frac{r^M - r^F}{SD^M}}_{\text{Market Sharpe Ratio}} SD^P \quad (10.9)$$

The left side of the equation is the expected return to efficient portfolios. The intercept is the risk-free return, r^F . The x -variable is the standard deviation of efficient portfolios, SD^P . The slope is the Sharpe ratio of the market portfolio, whose value in our example

is $(0.3324 - 0.031)/0.1622 = 1.86$. This value is very high and an artifact of the sample period used to estimate these values, January 2012 to December 2016, which was a period of particularly good stock market performance. The market Sharpe ratio from 1926 to 2021 is 0.44.

10.5.2 Security Market Line

The right plot in figure 10.7 presents the **security market line** or **SML**. The vertical axis measures expected returns just like that of the left plot. The key difference is the horizontal axis measuring risk. The right plot uses beta whereas the left plot uses standard deviation. The equation for the security market line was given earlier in equation 10.1 but is worth repeating.

$$r = r^F + \beta (r^M - r^F)$$

The security market line says that *any* asset's expected return equals the risk-free rate plus a risk premium that is determined by the asset's market risk exposure, β , and the market risk premium, $(r^M - r^F)$.

Comparing the security market line to the capital market line also reveals an important difference. If the CAPM is true (play along), then *all* assets lie on the security market line, whereas the only assets that lie on the capital market line are efficient portfolios (i.e., portfolios including the risk-free asset and the market portfolio). Take Boeing stock (BA) for example. BA lies in the interior of the mean-variance frontier for risky assets in the left plot because it is not efficient; there are other assets that offer the same return for less risk or more return for the same risk.

However, BA lies on the security market line in the right plot. Boeing has a beta, 0.582, and therefore we can use the SML to determine its expected return. We can estimate Boeing's beta from historical data but, in fact, we can back into Boeing's beta using a previous result. We know that Boeing and Efficient Portfolio 1 from the left plot have the same expected return. According to the CAPM, if two assets have the same expected return, then they must have the same beta.

Efficient portfolio 1 is a particularly simple portfolio; it contains the 41.8% invested in the risk-free asset and 58.2% invested in the market. Using equation 10.6, the beta of this portfolio is

$$0.418 \times 0 + 0.582 \times 1 = 0.582.$$

As we discussed earlier, the beta of a risk-free asset is zero, and the beta of the market portfolio is one. So, Boeing's beta is 0.582, the weight on the market in the efficient portfolio.

In practice, we would not rely on this approach. Rather, we would estimate Boeing's beta using historical data as we did with Apple above. Nonetheless, it is instructive to see the connection between the capital market line and the security market line. Both tell us the expected return to an asset. However, the capital market line tells us the expected returns for a very particular set of assets, namely, efficient portfolio consisting of the market portfolio and the risk-free asset while the security market line tells us the expected returns for any asset. And, the capital market line measures risk using volatility while the security market line measures risk using beta.

10.6 Risk-Adjusted Investment Performance

We know by now that when assessing the performance of our investments, we must take into account the risk of the investment. Riskier investments will generate higher returns, on average, just by their nature of being risky. One way that we've seen to adjust investments for risk is the Sharpe ratio, which divides the excess return - investment return less the risk-free return - on an investment by its volatility or standard deviation. The CAPM provides us with another measure of investment performance.

When we estimated Apple's beta, we took the covariance of its returns with the market returns and divided that number by the variance of the market returns. That is, the beta for stock "i" is

$$\beta^i = \frac{Cov^{i,M}}{Var^M}.$$

Another way to estimate beta results in a similar estimate. It begins with equation 10.4, which is repeated here and is referred to as a **population regression function**.

$$r^i = r^f + \beta^i (r^M - r^f) + e$$

Let's subtract the risk-free rate from both sides of the equation and add an intercept term, α , the Greek letter alpha.

$$\underbrace{(r^i - r^f)}_y = \alpha + \beta^i \underbrace{(r^M - r^f)}_x + e \tag{10.10}$$

This too is a population regression function. The y -variable is the excess return on asset i , $(r^i - r^f)$. The x -variable is the excess return on the market or market risk premium, $(r^M - r^f)$.

If we take expectations of both sides of equation 10.10, the idiosyncratic error term, e , drops out because it equals zero on average.

$$\underbrace{(r^i - r^f)}_y = \alpha + \beta^i \underbrace{(r^M - r^f)}_x$$

According to the CAPM, the only thing that should matter for expected returns is market risk exposure - the second term in the expression. Therefore, the intercept, α , should equal zero. If alpha is not zero, then the expected return is a function of something other than market risk.

We can estimate the parameters of this regression function - α and β - using **ordinary least squares** and the same returns data we used above to compute Apple's beta, January 2017 to December 2021. The only difference is rather than using raw returns for Apple and the market, we are using excess returns, i.e., subtracting off the risk-free return which we proxy for with the 30-day T-bill rate.

The data and estimated regression line are illustrated in figure 10.8. On the horizontal axis is our x -variable, the market excess return $(r_t^M - r_t^F)$. On the vertical axis is our y -variable, Apple's excess return $(r_t^{Apple} - r_t^F)$. There are 60 blue circles in the graph, each representing a pair of monthly excess returns between January 2017 and December 2021. For example, the data point towards the bottom of the plot labeled "2018:11" corresponds to November 2018. In that month, the excess return on the market was 1.7%; the excess return on Apple was -18.3%.

The solid black line is the estimated regression line. The vertical distance from each data point to the regression line is the estimated error, e , or **residual**. For example, on the left side of the plot, the residual for March of 2020 is indicated by the vertical bracket. So, we can think of every return as the sum of two components. The first is the value predicted by the CAPM, which is represented by the black line. The second is the error, which is represented by the vertical distance to the black line. Mathematically,

$$r_t^{Apple} - r_t^F = \underbrace{\hat{\beta}^{Apple} (r_t^M - r_t^F)}_{\text{Expected value}} + \underbrace{\hat{e}_t}_{\text{Residual value}}$$

The little "hat's" above β and e denote that these are estimates from our sample of returns. In our sample, the estimated slope of the regression line, $\hat{\beta}$, is 1.13. This is Apple's




Figure 10.8: Apple and Market Excess Returns and Estimated Regression Function

estimated beta, exactly what we estimated above.¹⁰

Take the data point denoted 2018:11 in Figure 10.8. The expected or predicted value for that month was

$$\hat{\beta}^{Apple} (r_{2018:11}^M - r_{2018:11}^F) = 1.13 \times 1.7\% = 1.9\%.$$

According to the CAPM, Apple stock should have returned 1.9% above the risk-free return for the month. But, it didn't, it returned -18.3% less than the risk-free return for the month. So, the residual or estimated error for November of 2018 was

$$\hat{e}_{2018:11} = (r_{2018:11}^{Apple} - r_{2018:11}^F) - \hat{\beta}^{Apple} (r_{2018:11}^M - r_{2018:11}^F) = -18.3 - 1.9 = -20.2\%.$$

That's a really big error for one month (or even a year!).

But, this big error does not mean that the CAPM is necessarily wrong or that Apple is a bad investment. In any given month (or day, week, year,...), the CAPM prediction will almost surely be incorrect because idiosyncratic events affect the stock returns every period. In other words, the primary prediction of the CAPM isn't that all the dots in Figure 10.8

¹⁰The estimate of beta obtained from the excess return regression here and that obtained earlier are slightly different, a fact hidden by rounding.

lie on the line. The primary prediction is that the CAPM is right *on average*; alpha equals zero and the line runs through the origin.

The y-intercept indicated by the red square in Figure 10.8 is the estimated alpha, which equals 1.95% per month - we're working with monthly returns - and is highly statistically significant. Annualized, the estimated alpha is $12 \times 1.95 = 23.4\%$. This quantity is called **Jensen's alpha**, and it has an important role in capital allocation and asset management. Jensen's alpha represents **abnormal return**. It is the component of expected returns not explained by risk or, more precisely in this case, market risk. Apple's alpha of 1.95% suggests it was a very good investment between 2017 and 2021. Apple's average monthly return over this period was 3.52% (42.28% per annum). Apple's expected monthly return over this period, according to the CAPM, was 1.57% (18.86% per annum). In other words, Apple was earning 1.95% per month (23.43% per annum) above and beyond what it should have earned given its market risk exposure or beta.

There are several potential explanations for Apple's positive alpha or outperformance. First, the CAPM is the wrong model. There are factors other than the market risk factor that are responsible for variation in asset returns, and alpha is capturing the effect of these omitted risk factors. We'll explore this possibility below. Second, our estimates of the market risk premium and risk-free rate are measured with error. This error biases our estimates of alpha (and beta) so that the true alpha is actually zero. Third, Apple was underpriced. In other words, the market was too pessimistic about Apple's prospects (underestimated future cash flows) or overestimated the future risk (i.e., discount rate). As a result of its low price, Apple has generated higher than expected returns according to the CAPM.

The problem with trying to take advantage of this second possibility is that we don't know if the market is wrong or our model is wrong. And, even if we know the market is wrong, we don't know *when* the market will figure out its mistake and correct it. In other words, let's say Apple stock is undervalued. We could buy the stock in the expectation that the price will increase and generate abnormally high returns. But, when will the price increase? In other words, the market has to agree with us at some point so that the price on Apple actually does increase. Related, perhaps the historical outperformance was an anomaly and the stock price today is correct or even overvalued. It's not obvious how to determine if the stock is over-, under-, or correctly-valued based on this analysis.

10.7 Other Factor Models

There are many (many, many,...) other asset pricing models, some used in practice by asset managers looking for an edge when investing their clients' or their own money and some used for performance assessment. These other asset pricing models arise primarily because of the empirical shortcomings of the CAPM, which doesn't seem to do a very good job of explaining expected returns. (Our Apple analysis is one example.) However, these alternative models have not gotten much use in estimating the cost of capital. Why? They often produce strange estimates (e.g. abnormally low or high and even negative costs of capital).

Let's take a look at one of the more popular alternatives to the CAPM, the Fama-French 3-factor model.¹¹ This model says that the expected returns for asset i are a function of three factors.

$$r^i = r^F + \underbrace{\beta^M (r^M - r^F)}_{\text{Market factor}} + \underbrace{\beta^{SMB} (r^{Small} - r^{Big})}_{\text{Size factor}} + \underbrace{\beta^{HML} (r^{High} - r^{Low})}_{\text{Value-growth factor}} \quad (10.11)$$

This model starts with the one-factor CAPM, equation 10.1, and adds two more risk factors.

1. **Size factor.** This factor was motivated by the observation that small firms, as measured by market capitalization, earned higher returns on average than big firms after accounting for differences in market betas across the two groups. In other words, if we took two firms - one large the other small - with the same market betas the smaller firm would have a higher average return. Thus, some researchers have argued that small firms must be riskier than big firms.

The size factor, often referred to as **SMB** for "small minus big," is computed as the difference between the returns to portfolios of small firms and large firms. β^{SMB} measures an asset's sensitivity to this risk factor. The estimated annual size risk premium, $r^{Small} - r^{Big}$, based on data from 1926 to 2021 is 2.37%, meaning small firms earned on average 2.37% higher returns than large firms over this period.

2. **Value-growth factor.** This factor was motivated by the observation that firms with high book-to-market equity ratios - so-called value firms - earned higher returns on average than firms with low book-to-market equity ratios - so-called growth firms -

¹¹This model is based on the findings in Eugene Fama and Kenneth French, 1992, "The Cross-Section of Expected Returns," *Journal of Finance* 47, 327-465.

after accounting for differences in market betas *and* firm size across the two groups. Thus, researchers have argued that value firms must be riskier than growth firms.

Book-to-market-equity is measured by the ratio of the book value of equity from the balance sheet to the market capitalization of the firm. The value-growth factor, often referred to as **HML** for “high minus low,” is computed as the difference between the returns to portfolios of high and low book-to-market equity firms. β^{HML} measures an asset’s sensitivity to this risk factor. The estimated annual value-growth risk premium, $(r^{High} - r^{Low})$, based on data from 1926 to 2021 is 4.04%, meaning value firms earned on average 4.04% higher returns than growth firms over this period.

The intuition behind the Fama-French 3-factor model is simple. The CAPM, relying only on market risk exposure, does a poor job of explaining variation in expected returns because we can find many assets with significant alpha - positive and negative. So, there must be other risks in the economy to which investors are exposed that affect expected returns and that are not accurately represented by market risk. SMB and HML do a good job, empirically, of proxying for those risks in that equation 10.11 eliminates alpha for many assets (but far from all). That is, the intercept in equation 10.11 is statistically indistinguishable from zero for many assets once we incorporate the size and value-growth factors into our expected return model.

10.7.1 Apple’s Equity Cost of Capital Revisited

Let’s estimate Apple’s equity cost of capital this time using the Fama-French 3-factor model.

Step 1: Estimate the Risk-free Rate

This step is unchanged from earlier. We’ll use the yield on the 30-year Treasury bond as of December 31, 2021, 1.90%, as a proxy for the risk-free rate.

Estimate Risk Premia

This step requires us to estimate not just the market risk premium, $(r^M - r^F)$, but also the size and value-growth risk premia. From earlier we estimated the market risk premium as 6% per annum. Using data from Ken French’s data library, estimates for the size and value-growth premia are 2.37% and 4.04%.

Estimate Betas ($\beta^M, \beta^{SMB}, \beta^{HML}$)

To estimate the betas in equation 10.11, we need to estimate a multivariate regression using historical data. Specifically, we estimate the following regression using monthly data from January 2017 to December 2021.

$$(r_t^{Apple} - r_t^F) = \alpha + \beta^M(r_t^M - r_t^F) + \beta^{SMB}(r_t^{Small} - r_t^{Big}) + \beta^{HML}(r_t^{High} - r_t^{Low}) + e_t$$

The coefficient estimates and corresponding t-values are as follows. The market beta, 1.33,

Coefficient	Estimate	T-value
α	1.13	1.30
β^M	1.33	6.94
β^{SMB}	-0.36	-1.12
β^{HML}	-0.76	-3.25

Table 5: Apple Alpha and Risk-Exposures in Fama-French 3-Factor Model

is similar to the CAPM estimate of 1.13, suggesting that Apple is somewhat riskier than investing in the market portfolio.

The coefficient on the size factor, -0.36, is statistically insignificant or statistically indistinguishable from zero. Taking the estimate at face value suggests that Apple has a negative exposure to the size factor. In other words, as small stocks become riskier relative to large stocks and earn higher returns, Apple's expected return should decline. This result is to be expected. Apple is a huge company - the second largest in the world as of late 2021 - with a market capitalization as of December 31, 2021 equal to \$2.91 trillion. It's returns are going to correlate positively with big firms whose return is subtracted in the size factor; hence, the negative coefficient.

The coefficient on the value-growth factor, -0.76, is highly statistically significantly different from zero. Like size, Apple's returns exhibit a negative risk exposure to the value growth factor, suggesting that despite its size, Apple's returns behave more like those of growth stocks than value stocks. As such, its expected return should be lower. (Remember, because growth stocks earn, on average, lower returns than value stocks they are viewed as less risky.)

Putting it all together

With our beta and risk premia estimates, we can compute the expected return to Apple using equation 10.11.

$$E(r^{Apple}) = 0.019 + \underbrace{(1.33 \times 0.06)}_{\text{Market Risk Premium}} - \underbrace{(0.36 \times 0.0237)}_{\text{Size Risk Premium}} - \underbrace{(0.76 \times 0.0404)}_{\text{Value-Growth Risk Premium}} = 0.0595$$

Apple's equity cost of capital is 5.95% according to the Fama-French 3-factor model, quite a bit lower than that implied by the CAPM (8.68%).

Is 5.95% crazy? Not really, especially as of December 31, 2021 when many market participants believed asset prices were inflated and future returns would be significantly lower going forward. That said, I would be shocked if Apple itself thought of its equity cost of capital being only 5.95%, and I am positive that Apple's shareholders are expecting a return greater than 5.95% per year. (I'm one of those shareholders:)

10.7.2 Parting Thoughts

Fama and French popularized two additional factors back in 1992. Since then there have been by some accounts hundreds of other factors thought to be responsible for expected returns. Why so many? Lots of reasons, though data-mining or p-hacking, and errors are popular explanations for many of them. Data mining and p-hacking is simply running a sufficiently large number of statistical tests, without appropriately accounting for the tests, until a statistically significant result is found.¹² In light of this factor chaos, it's not surprising why the CAPM still reigns supreme for computing the cost of capital in practice. It's simple, easy to apply, and not terribly controversial even if it has its empirical failings.

10.8 Key Ideas

The CAPM makes some strong (unrealistic) assumptions but generates some elegant and useful predictions.

¹²Campbell Harvey and Yan Liu document over 400 factors found by academics in their study: Harvey and Liu, 2020, "A Census of the Factor Zoo," Working Paper, Duke University. Studies by Linnainmaa and Roberts, 2018, History of the cross-section of stock returns, *Review of Financial Studies* 31, 2606–2649; and Chordia, Goyal, and Saretto, 2020, Anomalies and false rejections, *Review of Financial Studies* 33, 2134–2179 present evidence of data-mining and p-hacking.

- The tangency portfolio that identifies the efficient portfolio of risk assets is the market portfolio - a value-weighted portfolio containing all assets in the economy. We often proxy for this portfolio with broad-based stock market index funds, like Vanguard's total stock market fund or S&P 500 fund.
- All investors should hold some combination of the risk-free asset (e.g., Treasury security) and the market portfolio (e.g., Total stock market index fund or S&P500 tracker fund) because together they form mean-variance efficient portfolios.
- The capital market line relates the portfolio volatility to expected returns for portfolios containing the risk-free asset and the market portfolio.

$$r^P = r^F + \underbrace{\frac{r^M - r^F}{SD^M}}_{\text{Market Sharpe ratio}} \times SD^P$$

- Total risk is comprised of two components: (i) Systematic (market, non-diversifiable, priced) risk and (ii) Idiosyncratic (diversifiable firm-specific, non-priced) risk. The former reflects risks that affect all assets; the latter affects a subsets of or individual assets. Because the latter doesn't bring any benefit in terms of more return, investors should seek to eliminate all idiosyncratic risk by investing in diversified portfolios - according to the CAPM.
- The security market line relates beta, our measure of market risk exposure, to expected returns for all assets.

$$r^i = r^F + \beta^i (r^M - r^F)$$

- Jensen's alpha measures abnormal returns that are in excess (when positive) or dearth (when negative) of what is expected given the risk exposure of the asset. Investors like positive alpha because it suggests they are getting more return than they should given their risk exposure.
- There are many other models of expected returns (e.g., Fama-French 3-factor model), whose popularity is concentrated in the asset management space, as opposed to corporate finance. Investors use these models to identify alpha, but companies and banks have yet to embrace them to estimate cost of capitals.

10.9 Technical Appendix

10.9.1 Adjusting Beta Estimates

Because estimates of beta can be quite noisy, several adjustments have been proposed. The **Bloomberg adjustment** attempts to mitigate large magnitude beta estimates with the following adjustment.

$$\text{Bloomberg beta} = 0.66 + 0.34 \times \beta \quad (10.12)$$

This adjustment shrinks betas towards one. So, betas larger than one are reduced. Betas smaller than one are increased.

Another adjustment suggested in the study of Handa, Kothari, and Wasley addresses the problem of small stocks reacting to market news with a lag.¹³ Though this problem is mostly avoided by using monthly, as opposed to daily or weekly, return data to estimate beta, these authors suggest incorporating a lagged market return in the regression specification to estimate beta like so.

$$(r_t^i - r_t^F) = \alpha + \beta_1(r_t^M - r_t^F) + \beta_2(r_{t-1}^M - r_{t-1}^F) + e_t$$

The adjusted beta is the sum of the β_1 and β_2 estimates.

10.10 Problems

10.1 (*Conceptual*) Determine whether each of the following statements are true or false according to the CAPM.

- The capital asset pricing model (CAPM) is an accurate representation of reality.
- All investors should hold a portfolio containing the risk-free asset and the market portfolio.
- All investors share the same expectations about future asset prices.
- The skewness and kurtosis an asset's return distribution is important to investors.
- Risk averse investors should hold portfolios containing more of the risk-free asset relative to risk-tolerant investors.

¹³Handa, P, S.P. Kothari, and C. Wasley, 1989, "The relation between the return interval and betas: Implications for the size effect," *Journal of Financial Economics* 23, 79-100.

- f. On idiosyncratic or firm-specific risk matters to investors and therefore affects expected returns.
- g. The volatility of a firm's stock returns determines the its expected return
- h. Every asset in the economy lies on the security market line (SML)
- i. Assets with a beta of zero have an expected return equal to the risk-free rate.
- j. Firms with higher firm-specific risk should garner higher expected returns.

10.2 (*Risk premium, expected return, expected price*) MillerKnoll, Inc. (MLKN) is a furniture designer, manufacturer, and distributor. As of January 2024, MLKN's stock price was \$26.49 per share and their beta was 1.32. The 1-year yield on Treasury bills was 4.80% and the market risk premium was 5%.

Using this information, answer the following questions.

- a. What the risk premium on MLKN stock?
- b. What is MLKN's expected return according to the CAPM?
- c. What MLKN's expected share price one year from today?
- d. How does the market risk exposure of MLKN compare to that of the well diversified portfolio?

10.3 (*Relative mispricing, alpha*) As of January 2024, Rivian automotive (RIVN) was trading for \$19.91 a share with a beta of 2.13, Livent Corporation (LTHM) was trading for \$16.51 a share with a beta of 1.76. The market risk premium is 5% and the yield on a one-year Treasury bill is 4.68%.

Using this information, answer the following questions.

- a. What are the expected returns for Rivian and Livent according to the CAPM?
- b. What are the expected share prices for Rivian and Livent one year from today?
- c. Analyst forecasts suggest that Rivian's expected return is 21%. If analysts are correct, is Rivian over- or under-priced in the market and should investors buy or (short-) sell Rivian in January 2024? What is Rivian's alpha?
- d. Analyst forecasts suggest that Livent's expected return is 9.5%. If analysts are correct, is Livent over- or under-priced in the market and should investors buy or (short-) sell Livent in January 2024? What is Livent's alpha?

- e. Draw the security market line (SML) and plot Rivian's and Livent's betas and expected returns. What is the relation between these stocks' estimates and the SML?
- 10.4 (*Zero-beta assets, risk-free returns*) In December 2022, TAL Education Group (TAL) had a beta of 0. If the one-year Treasury yield was 4.25% and the market risk-premium 6.5%, what is the expected return to TAL? How does it compare to the risk-free return? Explain.
- 10.5 (*Portfolio betas*) Consider a portfolio consisting of only a Treasury bill and the market portfolio. Answer the following questions.
- What is the beta of a portfolio consisting of 25% T-bill and 75% market?
 - What is the beta of a portfolio consisting of 50% T-bill and 50% market?
 - If we want a portfolio beta of 1.5, what must the portfolio weights be?
 - What is the general relation between the weight on the market portfolio and the beta of the portfolio?
- 10.6 (*Market neutral portfolios and returns*) Using the two stocks in the following table, construct a market neutral portfolio. If the market risk premium is 5% and the risk-free return is 3.5%, what is the expected return on the market neutral portfolio you constructed?
- | Stock | Beta |
|-----------------------|------|
| Nextera Energy (NEE) | 0.52 |
| JP Morgan Chase (JPM) | 1.11 |
- 10.7 (*Mutual fund performance*) Saul Katz manages a mutual fund that is benchmarked to the S&P 500 index. In other words, Saul's performance is measured each year by comparing the return to his portfolio to that of the S&P 500. In 2013, the return to the S&P 500 was 29.6% and the yield on a one-year Treasury was 0.15%. That same year, Saul's fund returned 31.4%. Based on just this information, how did Saul do? If you also knew that Saul's portfolio beta was 1.25, how would you gauge Saul's performance? Explain any difference between the two scenarios.
- 10.8 (*Stock valuations*) Vaxcyte Inc. (PCVX) is a biotech vaccine company with a market beta of 0.92. Riot Platforms Inc. (RIOT) is a bitcoin mining company with a market beta of 4.05. The market risk premium is 6% and the yield on a one-year Treasury bill is 5.2%. Analysts forecast a 14% return for Vaxcyte and 22% return for Riot. Using this information, answer the following questions.

- a. What are the expected returns to Vaxcyte and Riot?
- b. Based on the analyst forecasts, are Vaxcyte and Riot over-, under-, or correctly-valued?

10.9 (*Alpha and beta estimation*) SunRun Inc. is designer and installer of solar panels in residential real estate. Their weekly stocks returns, along with the returns to the S&P500 and the three-month T-bill, are presented in the table below. (All returns are in percentages.)

Date	SunRun (RUN)	S&P 500	T-bill
1/5/24	(13.09)	0.45	0.02
1/1/24	1.92	(1.70)	0.08
12/25/23	3.38	0.32	0.55
12/18/23	44.53	0.75	(0.35)
12/11/23	(10.11)	2.49	0.11

Using this data, answer the following questions.

- a. What are the average weekly returns to each of the three assets? What are the corresponding annualized average returns?
- b. What is SunRun's beta?
- c. What is SunRun's alpha in both weekly and annual terms?
- d. If the risk-free rate is 4% and the market risk premium 5%, what is SunRun's weekly expected return according to the CAPM? What is its annual expected return according to the CAPM?
- e. How reliable do you think your estimates are? Explain.

10.10 (*Beta, covariance, and correlation*) Fill in the question marks in the table below.

	Volatility	Covariance with Market	Beta	Expected Return (%)
Stock A	0.40	0.20	?	?
Stock B	0.60	?	1.80	14.60
Risk-free asset	?	?	?	6.50
Market	0.25	?	?	?

10.11 (*Risk decomposition*) Constellation Brands, Inc. (STZ) is a beverage company whose annual stock return information based on monthly returns from January 2017 and December 2022 is presented in the table. Also presented in the table is corresponding information for the broader market.

	STZ	Market
Average (%)	11.39	12.14
Volatility (%)	24.83	16.87
Beta	0.94	1.00

Using this information, answer the following questions.

- If the risk-free rate is 4% per annum, what is the expected return to STZ according to the CAPM? How does this compare to STZ's realized average return?
 - What is the total risk of STZ stock? (Express your answer in terms of Variance.)
 - What is the systematic risk of STZ stock? (Express your answer in terms of Variance.)
 - What is the idiosyncratic risk of STZ stock? (Express your answer in terms of Variance.)
 - What proportion of return variation in STZ stock comes from market risk versus firm-specific risk?
- 10.12 (*Equity cost of capital*) Assuming the yield on the 30-year Treasury is 6.2% and the market risk premium is 5%, estimate the equity cost of capital for each firm in the table below. Also provide a brief explanation for why each stock has a higher or lower expected return.

Stock	Beta	Description
Unity Software (U)	2.53	Platform for software development tools and services
Hewlett Packard Enterprises (HPE)	1.26	Data solutions and computer hardware retailer
The Gap (GAP)	2.17	Clothing designer and retailer
Novavax Inc (NVAX)	1.52	Biotech company focused on vaccines
Duke Energy Corp. (DUK)	0.47	Utility company

- 10.13 (*Beta interpretation*) Kiniksa Pharmaceuticals LTD (KNSA) had a market beta equal 0.00 as of December 2022. The market risk premium was 5%, and the risk-free rate was 4.12%. Using this information and assuming the assumptions of the CAPM hold, answer the following questions.

- What is KNSA's expected return?

- b. What is the risk premium on KNSA stock?
- c. If the annual volatility of KNSA stock is 29%, is there an arbitrage opportunity?

10.14 (*Efficient portfolios*) Callie Hartman recently received a portfolio of stocks from the estate of her late grandfather. The portfolio has market beta of 0.6 and annual volatility of 38%. A broad market index has an annual expected return and volatility of 10% and 20%, respectively. The risk-free return is 5.4%.

Using this information, answer the following questions.

- a. What is the market risk premium?
- b. What is the expected return to Callie's portfolio?
- c. What is the Sharpe ratio of Callie's portfolio?
- d. What is the Sharpe ratio of the market?
- e. Is Callie's portfolio mean-variance efficient?
- f. Find the portfolio with the same expected return as Callie's portfolio and with the lowest volatility? What assets are in this portfolio? What are the portfolio weights? What is the beta of this portfolio? What is the volatility of this portfolio? What fraction of Callie's portfolio's volatility is systematic and what fraction is idiosyncratic?
- g. Find the portfolio with the same volatility as Callie's portfolio and with the highest expected return. What assets are in this portfolio? What are the portfolio weights? What is the beta of this portfolio? What is the expected return of this portfolio?

Chapter 11

Optimal Financial Policy

Fundamental value relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

We explored financing way back in chapter 3 when we examined the implications of borrowing money for personal decision making. What we didn't do is ask: How *should* we finance a purchase? Put differently, what are the implications of different financing choices for value and rates of return?

This chapter

- provides a conceptual framework to understand the effects of different financial policies on asset returns and values,
- shows how to compute the cost of capital for projects with any financing scheme,
- debunks and clarifies common fallacies surrounding corporate financing,
- applies these lessons to answer several questions including:
 - What is Apple Inc.'s debt and equity costs of capital? What is Apple Inc.'s weighted average cost of capital?
 - What impact does a leveraged recapitalization have on a company's debt, equity, and weighted average costs of capital?

- What impact does payout policy - dividends and share repurchases - have on a company’s debt, equity, and weighted average costs of capital?

These questions all relate to financial policy and **capital structure**, which is the mix of financial instruments used to fund an asset or collection of assets. Financial policy is the most important part of a chief financial officer’s (CFO’s) job according to survey evidence.¹ It is also one of the most difficult. We’ll start by exploring how financing choice affects outcomes under a very strict, and unrealistic, set of assumptions called **perfect capital markets**. We’ll then relax assumptions to show how financial policy can create value.

11.1 Market Value Balance Sheet

Let’s start with a really useful construct - the **market value balance sheet** or **MVBS**. A MVBS is just like a regular accounting or “book” balance sheet with assets on the left side and claims - liabilities and shareholders equity - on the right side. If we assume there are no arbitrage opportunities, an MVBS also has to balance: total assets = total claims. The difference between book and market value balance sheets is the values assigned to each account. Values in an accounting balance sheet are based on historical transactions, values in a market value balance sheet are based on discounted future cash flows. In other words, the value of each account in an MVBS follows from our fundamental value relation.

Table 1 presents the market value balance sheet for Granite Construction Inc. (GVA), an infrastructure construction and mining company, as of May 15, 2024. The claims on the firm are separated into debt and equity. If we wanted greater resolution into the company, we could break out each of these categories into their components. For example, we could have listed all four of GVA’s debt instruments - a revolving line of credit, a term loan, and two convertible notes, as well as their lease agreements. Likewise, we could have listed both preferred and common stock, though GVA doesn’t have any preferred stock. We also could have broken out the operating assets of GVA into their two operating segments: construction and materials. The decision to do so depends upon our goals. Regardless, the mechanics and intuition we discuss below are the same so we’ll focus on the level of aggregation presented in Table 1 to avoid unnecessary complexity.

Focusing first on the claims of the MVBS, the equity value is the market capitalization of GVA’s common stock. The debt value is the sum of outstanding debt and leases as reported

¹Survey evidence compiled in Henri Servaes and Peter Tufano, 2006, “CFO Views on the Importance and Execution of the Finance Function,” Deutsche Bank shows that CFOs believe capital structure and debt issuance and management are the two finance functions that add the most value to a company.

on the balance sheet of their most recent quarterly 10-Q filing with the Securities Exchange Commission (SEC). Summing the claim values - debt and equity - gives us the value of the firm. It is worth emphasizing that the value of a firm is not the same as the market capitalization unless the firm only has common stock and no other financing arrangements (e.g., debt, leases, preferred stock).

Assets		Claims	
Cash, C	506.24	Debt, D	633.57
Operating, O	2,886.33	Equity, E	2,759.00
Total assets, V	3,392.57	Firm Value, V	3,392.57

Table 1: Market Value Balance Sheet - Granite Construction Inc. (GVA) May 15, 2024. Values in \$millions

Debt values from the accounting balance sheet - so-called **book values** - are just an approximation of the market values of debt. This approximation will be poor if either interest rates or the financial health of the firm have changed significantly since the debt was originally issued. The book value will be larger than the market value if interest rates have increased or the financial health of the firm deteriorated. The book value will be smaller than the market value of debt if interest rates have decreased or the financial health of the firm improved. Estimating market values in response to changes in interest rates or financial health can be accomplished for fixed rate bonds using the examples discussed in chapter 7. However, some types of debt - e.g., floating-rate, convertible, callable - require tools beyond the scope of this text. To value these instruments, see the excellent text by Pietro Veronesi, "Fixed Income Securities: Valuation, Risk, and Risk Management."

Turning to the assets of the MVBS, we start with cash which can be obtained from the accounting balance sheet. Cash refers to all cash, cash equivalents, and short-term investments found in the current assets section of the book balance sheet. Because these are all highly liquid and short-term assets, if not literally cash, the value reported on the accounting balance sheet often offers a reasonable approximation to the market value. This approximation will be poor if the firm has used or added to its cash between the date of valuation and the last balance sheet reporting date. We'd need to search recent announcements by the firm or filings (8K forms) to determine if such a change has occurred.

The other category of assets are operating assets, which refer to all other assets of the firm. This category encompasses the assets found on the accounting balance sheet - plant, property, equipment, intangible assets, and financial assets - as well as assets not on the accounting balance sheet - brand, human capital, and customer loyalty. However, the values

of these assets corresponds to market values or the sum of the present values of their future cash flows.

11.1.1 Enterprise Value and Excess Assets

We can estimate the value of operating assets two ways. First, the sum of cash and operating assets must equal total assets implying operating assets equal total assets minus cash.

$$\text{Total assets} = \text{Operating assets} + \text{Cash} \implies \text{Operating assets} = \text{Total assets} - \text{Cash}$$

The value of GVA's operating assets is $3,392.57 - 506.24 = \$2,886.33$ million. Alternatively, because the balance sheet must balance, operating assets equals equity plus debt minus cash.

$$\text{Cash} + \text{Operating assets} = \text{Debt} + \text{Equity} \implies \text{Operating assets} = \text{Equity} + \text{Debt} - \text{Cash}$$

Again, the value of operating assets is $2,759.00 + 633.57 - 506.24 = \$2,886.33$ million.

The value of operating assets is more commonly referred to as **enterprise value**, and represents the value of the business net of any **excess assets**. An excess asset is an asset that the firm does not need to run the business. For example, a company may have a defunct production facility, old computers, unused office furniture, etc. These assets may generate cash flows and have value that can be realized on secondary markets, but they don't contribute to the generation of cash flows essential to the business. They can be liquidated and sold without impacting the free cash of the business.

Because accounting for and valuing all excess assets is exceedingly difficult, practitioners treat all cash on the balance sheet as an excess asset, hence the taxonomy in Table 1. While convenient, assuming all cash is excess is clearly unrealistic. No firm can operate without cash, which is needed to cover short-term liabilities like payroll, utilities, and rent. In other words, part, if not all, of the cash on the balance sheet is **required cash** and a part of working capital. The relevant question is: What fractions of the cash are required versus excess? This is a difficult question to answer especially for people outside of the company and probably why the convention of treating all cash on the balance sheet like an excess asset persists. Nonetheless, it is important to understand this assumption so that we can modify it when necessary.

To summarize, the MVBS is useful visual for organizing the market values of a firm's assets and claims. We relied on GVA being a publicly traded company to complete its market value balance sheet. For privately held firms, we need to undertake some additional steps, which we'll explore in chapter 12.

11.1.2 Expected Returns and Beta

Now let's use the MVBS to uncover the relations between the expected returns to assets and claims. Because the CAPM plays a central role in estimate expected returns, let's also explore the relations among the market betas. Table 2 presents the market value balance sheet with expected returns and market betas in place of the dollar values. The table construction began by estimating the equity cost of capital using the CAPM. GVA's equity beta from Yahoo! Finance was 1.46. With a 4.43% yield on a 10-year Treasury note and a 5% market risk premium, the estimated equity cost of capital is $4.43 + 1.46 \times 5.00 = 11.73\%$. In other words, shareholders expect to earn 11.73% per year or, equivalently, it costs GVA 11.73% per year to raise equity capital.

Assets		Claims	
Expected return on cash, r^C (%)	7.46	Debt cost of capital, r^D (%)	7.46
Cash beta, β^C	0.61	Debt beta, β^D	0.61
Expected return on operating assets, r^O (%)	11.54	Equity cost of capital, r^E (%)	11.73
Operating beta, β^O	1.42	Equity beta, β^E	1.46
Expected return on assets, r^A (%)	10.93	Expected return on assets, r^A (%)	10.93
Asset beta, β^A	1.30	Asset beta, β^A	1.30

Table 2: Market Value Balance Sheet - Expected Returns and Market Betas - Granite Construction Inc. (GVA) May 15, 2024. Values in \$millions

The debt cost of capital was estimated at 7.46% using information in GVA's most recent annual filing. Assuming the debt is trading close to par and the firm is not in financial distress, using the coupon rate or yield on the debt can provide a reasonable estimate of the expected return. For debt that is significantly impaired, we can employ the approach in the technical appendix of chapter 7. With the expected return to debt, we can back out the debt beta using the CAPM.

$$r = r^F + \beta \times (r^M - r^F) \implies \beta = \frac{r - r^F}{r^M - r^F} = \frac{7.46 - 4.43}{5.00} = 0.61$$

The claims cost of capital, which is just the asset cost of capital because the MVBS must balance, is a portfolio of debt and equity. To compute the return or beta, we take a weighted average of the component returns where the weights are component values divided by the

portfolio value.

$$r^A = \frac{Debt}{Debt + Equity} \times r^D + \frac{Equity}{Debt + Equity} \times r^E \quad (11.1)$$

$$\beta^A = \frac{Debt}{Debt + Equity} \times \beta^D + \frac{Equity}{Debt + Equity} \times \beta^E$$

Equation 11.1 is referred to as the **pre-tax weighted average cost of capital** or **pre-tax WACC**. It tells us the cost of raising money from *all* claimants. Using the values from Tables 1 and 2, we get

$$r^A = \frac{633.57}{633.57 + 2,759.00} \times 7.46 + \frac{2,759.00}{633.57 + 2,759.00} \times 11.73 = 10.93\%, \text{ and}$$

$$\beta^A = \frac{633.57}{633.57 + 2,759.00} \times 0.61 + \frac{2,759.00}{633.57 + 2,759.00} \times 1.46 = 1.30.$$

Turning to the assets, the expected return on cash is *assumed* to be the same as the debt cost of capital, 7.46%. This is a simplifying assumption that often avoids a logical inconsistency. Remember, cash is assumed to be an excess asset, otherwise it would be part of working capital and therefore operating assets. If the return on cash was less than debt, then it would be value-destructive for the firm to have any debt. This is akin to homeowners saving extra money in an investment vehicle with a lower return than their mortgage. In this situation, the homeowner, like the company, would be better off paying down the relatively high cost debt. If instead the cash was earning a higher return than debt, then the company would have a fantastic business opportunity because most any financial institution would gladly lend against cash collateral. Typically, most firms earn somewhat less on the excess cash than what they pay on their debt largely because they do not delineate between excess and required cash. The cash beta can be found in the same way the debt beta was found by using the CAPM.

The operating asset expected return can be estimated several ways. First, we can recognize that the assets of the firm are just a portfolio of cash and operating assets. Therefore, the return on assets is a weighted average of the return on cash and operating assets.

$$r^A = \frac{Cash}{Cash + Operating} \times r^C + \frac{Operating}{Cash + Operating} \times r^O$$

Solving for the operating asset expected return yields the following.

$$\begin{aligned} r^O &= \frac{Assets}{Operating} \times r^A - \frac{Cash}{Operating} \times r^C \\ &= \frac{3,392.57}{2,886.33} \times 0.1093 - \frac{506.24}{2,886.53} \times 0.0746 \\ &= 11.54\% \end{aligned}$$

Equivalently, the return to both sides of the balance sheet must be equal.

$$\begin{aligned} & \underbrace{\frac{Cash}{Cash + Operating} \times r^C + \frac{Operating}{Cash + Operating} \times r^O}_{\text{Return on assets, } r^A} \\ = & \underbrace{\frac{Debt}{Debt + Equity} \times r^D + \frac{Equity}{Debt + Equity} \times r^E}_{\text{Return on claims, } r^A} \end{aligned}$$

Solving for the operating asset expected return and assuming that the expected return on cash (r^C) is the same as the expected return on debt (r^D), we get the following expression.

$$r^O = \frac{Debt - Cash}{Operating} \times r^D + \frac{Equity}{Operating} \times r^E$$

Because we can more easily obtain values for the debt, equity, and cash, let's rewrite this equation like so.

$$r^O = \frac{Debt - Cash}{Equity + Debt - Cash} \times r^D + \frac{Equity}{Equity + Debt - Cash} \times r^E \quad (11.2)$$

Plugging values from Tables 1 and 2 into equation 11.2 produces the same result as before.

$$r^O = \frac{633.57 - 506.24}{2,759.00 + 633.57 - 506.24} \times 7.46 + \frac{2,759.00}{2,759.00 + 633.57 - 506.24} \times 11.73 = 11.54\%$$

Equation 11.2 expresses the business risk as a weighted average of the debt and equity cost of capitals where the weights are debt minus cash, i.e., **net debt**, and equity scaled by the enterprise value of the firm. We'll refer to net debt divided by enterprise value as **net leverage**, which is a measure of the net debt financing of the operations.

GVA's operating cost of capital is quite similar to its equity cost of capital because the company has almost as much cash as it does debt. Because cash and debt have the same expected return, cash acts like **negative debt**. Each dollar of cash is equivalent to one less dollar of debt and each dollar of debt is equivalent to one less dollar of cash. Thus, GVA's capital structure is not all that different from an **unlevered firm**, one with no excess cash and no debt, because its net leverage, $(633.57 - 506.25)/2,886.33 = 4.41\%$ is close to zero.

Betas for cash, operating assets, and the firm can be obtained using the estimated returns and the CAPM, or we can simply replace r in each equation with β . For example,

$$\beta^O = \frac{Debt - Cash}{Equity + Debt - Cash} \times \beta^D + \frac{Equity}{Equity + Debt - Cash} \times \beta^E \quad (11.3)$$

While we estimated the equity cost of capital, r^E , using the CAPM, we can also use equations 11.1 or 11.2 to express the equity cost of capital in terms of either the asset and debt cost of capitals or the operating and debt cost of capitals.

$$r^E = r^A + \frac{Debt}{Equity} (r^A - r^D) \quad (11.4)$$

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) \quad (11.5)$$

Equations 11.4 and 11.5 show that the equity cost of capital is equal to the asset or operating cost of capital plus an adjustment for financial risk. As debt financing increases, the equity cost of capital increases. Hence, the equity cost of capital is often referred to as the **levered cost of capital**. We'll see why in the next section.

11.1.3 Summary

Table 3 combines all of the results for Granite Construction - dollar values, expected returns/cost of capitals, and betas. The MVBS is not only a useful tool for organizing the values of assets and claims. It also illuminates the relationships between the expected returns and betas of assets and claims, both of which are central in capital budgeting - chapters 5 and 6 and corporate valuation 12. For example, the operating cost of capital, r^O , is the discount rate companies can use to evaluate projects with risks similar to their business or the free cash flows that comprise the company's operations. The asset cost of capital, r^A , is the discount rate to value the entire company - operating assets and any excess cash because it captures the risk of both. The two are the same when there is no excess cash. However, to estimate these discount rates, it is often easier to do so by first estimating the expected returns on debt, equity, and cash in conjunction with equations 11.1 and 11.2.

Assets		Claims	
Cash, C	506.24	Debt, D	633.57
Cash expected return, r^C (%)	7.46	Cost of debt, r^D (%)	7.46
Cash beta β^C	0.61	Debt beta, β^D	0.61
Operating, O	2,886.33	Equity, E	2,759.00
Operating expected return, r^O (%)	11.54	Cost of equity, r^E	11.73
Operating beta, β^O	1.42	Equity beta, β^E (%)	1.46
Total assets, V	3,392.57	Firm Value, V	3,392.57
Asset expected return, r^A (%)	10.93	Asset expected return, r^A (%)	10.93
Asset beta β^A	1.30	Asset beta β^A	1.30

Table 3: Market Value Balance Sheet - Granite Construction Inc. (GVA) May 15, 2024. Values in \$millions

Let's summarize some key concepts and relations to come from our discussion thus far.

- Firm value.

$$\text{Firm value} = \text{Equity} + \text{Debt} \quad (11.6)$$

Firm value is the value of *all* assets or, equally, all claims - debt and equity.

- Leverage.

$$\text{Leverage} = \frac{\text{Debt}}{\text{Equity} + \text{Debt}} \quad (11.7)$$

Leverage measures the fraction of firm or total asset financing that comes from debt.

- Asset cost of capital, r^A .

$$r^A = \frac{\text{Debt}}{\text{Debt} + \text{Equity}} \times r^D + \frac{\text{Equity}}{\text{Debt} + \text{Equity}} \times r^E$$

The asset cost of capital is the expected return to *all* of the firm's assets - operating and excess cash - or equally all of the firm's claims - debt and equity.

- Enterprise value.

$$\text{Enterprise value} = \text{Equity} + \text{Debt} - \text{Cash} = \text{Equity} + \text{Netdebt} \quad (11.8)$$

Enterprise value is the value of *operating* assets or business net of excess assets.

- Net leverage.

$$\text{Net leverage} = \frac{\text{Debt} - \text{Cash}}{\text{Equity} + \text{Debt} - \text{Cash}} \quad (11.9)$$

Net leverage measures the fraction of enterprise or operating asset financing that comes from debt less excess cash.

- Operating cost of capital, r^O .

$$r^O = \frac{\text{Debt} - \text{Cash}}{\text{Equity} + \text{Debt} - \text{Cash}} \times r^D + \frac{\text{Equity}}{\text{Equity} + \text{Debt} - \text{Cash}} \times r^E$$

The operating cost of capital is the expected return generated by the operations of a company. It is the appropriate discount rate for the operating free cash flows of the company and free cash flows of projects with risk similar to that of the company.

- Equity cost of capital, r^E .

$$\begin{aligned} r^E &= r^A + \frac{\text{Debt}}{\text{Equity}} (r^A - r^D) \\ r^E &= r^O + \frac{\text{Debt} - \text{Cash}}{\text{Equity}} (r^O - r^D) \end{aligned}$$

These equity cost of capital is the return expected by shareholders and the cost to the company from raising equity capital. It equals the asset (or operating) cost of capital plus a premium for any financing leverage which creates risk above and beyond that coming from the business.

11.2 Financing a Home

How does capital structure affect value and expected returns. In other words, as we vary the fraction of debt and equity in our market value balance sheet, what happens to the values of debt, equity and assets? What happens to the expected returns to debt, equity, and assets? Because finance principles are the same regardless of application, we'll answer these questions in a setting with which we're familiar - financing a home. What's useful about this example is that in addition to providing a setting with which we're already familiar, financing a home and financing a business are very similar. Just like a corporation finances assets with debt and equity, a house is an asset financed with debt (mortgage from a bank) and equity (cash from the homeowner). So, while we refer to a "house" and a "mortgage" below, we could change those labels to "equipment" and "loan," or "plant" and "bond" with no impact on the results and lessons.

11.2.1 Setup

Imagine we purchase a home for a \$1 million today that we'll sell one year from today. Let's also assume that there are three prices for which we could sell the home, each with equal probability: \$800,000, \$1.1 million, or \$1.3 million. We'll refer to these possibilities as the Low, Medium, and High states of the world. This information is summarized in table 4.

	Low	Medium	High
Future sales price (\$1,000)	800	1,100	1,300
Probability of outcome (%)	33.3	33.3	33.3

Table 4: Home Sale Price Scenarios and Probabilities

We'll abstract from some real-life details for two reasons. First, some of these details make the calculations more tedious but have no effect on the lessons we want to communicate. For example, let's assume the mortgage is an annual coupon bond instead of an amortizing bond. That is, the mortgage requires a fixed annual interest payment equal to the loan rate of 5% times the loan principal, which is due at maturity. Second, some of the details make the analysis more difficult, but also more interesting. We'll discuss these details further below. For now, let's think of the analysis of this section as establishing a baseline set of results on which we can build by incorporating more real life details.

11.2.2 The Capital Structure Choice

Table 5 presents house, debt, and equity values under different capital structures at the time of purchase. For example, the first row illustrates the case in which we pay all cash for the home, and the last row illustrates the case in which we pay \$50,000 and borrow the

remaining \$950,000. Each row of Table 5 is a market value balance sheet with zero excess cash ($C = 0$) and differing from one another only in how the asset is financed. The leverage column is the ratio of debt to assets.

Leverage (%)	Balance Sheets		
	Assets	Debt	Equity
0	1,000	0	1,000
25	1,000	250	750
50	1,000	500	500
75	1,000	750	250
95	1,000	950	50

Table 5: Debt and Equity Values (\$000s) for Different Capital Structures at Time of Purchase

11.2.3 Future Cash Flows and Values

Fast forward one year from today when we sell the home. Table 6 presents the cash flows generated by the house (i.e., the sale price) and how those cash flows are divided between the lender and the homeowner. Remember, in the low state we sell the home for \$800,000, in the medium state \$1.1 million, and in the high state \$1.3 million.

Leverage (%)	Debt						Equity			Asset		
	Interest			Principal			Low	Medium	High	Low	Medium	High
	Low	Medium	High	Low	Medium	High						
0	0.0	0.0	0.0	0	0	0	800	1,100	1,300	800	1,100	1,300
25	12.5	12.5	12.5	250	250	250	538	838	1,038	800	1,100	1,300
50	25.0	25.0	25.0	500	500	500	275	575	775	800	1,100	1,300
75	37.5	37.5	37.5	750	750	750	13	313	513	800	1,100	1,300
95	0.0	150.0	176.5	800	950	950	0	0	173	800	1,100	1,300

Table 6: Cash Flows (\$1,000s) at time of Sale for Different Capital Structures

Starting with the first row in which the home was entirely equity financed (no mortgage), there are no interest or principal payments. Therefore, the value of the equity - how much the homeowner receives - one year from today is equal to the sales price. Now consider the middle row in which leverage is 50%, meaning we borrowed \$500,000 and wrote a check for \$500,000 to purchase the house. At a 5% annual interest rate, we owe the bank $\$500,000 \times 0.05 = \$25,000$ in interest, and \$500,000 in principal repayment when we sell one year later. The cash flows are the same in the Low, Medium, and High states of the world because our payment to the bank is the same regardless of the sale price of the home. Our equity cash flow is what's left over after we've paid the bank. Remember, equity is the **residual claimant**, i.e., they only get paid if there is money left over after paying the creditors. For example, if we sell our home for \$800,000, our equity value is $800,000 - (25,000 + 500,000)$

= \$275,000. If we sell our home for \$1.3 million, our equity value is $1,300,000 - (25,000 + 500,000) = \$775,000$.

Finally, consider the last row in which we borrowed \$950,000. The interest expense on a 5% loan is $\$950,000 \times 0.05 = \$47,500$. In one year, we'll owe the bank a total of $47,500 + 950,000 = \$997,500$. If we sell our home in the low state of the world for \$800,000, we'll default on our loan. That is, we won't have enough money from the sale of our house to repay the bank - $\$800,000 < \$997,500$. Therefore, the loan is risky, and the bank is going to charge us a higher interest rate to compensate for that risk. How much higher? The technical appendix at the end of the chapter shows the bank requires an 18.58% interest rate when we borrow \$950,000 to compensate for the additional risk.² This higher interest rate means that with a \$950,000 loan, we'll default in two states of the world - Low and Medium - because the sales prices - \$800,000 and \$1,100,000 - are both lower than the amount we'll owe to the bank - $\$950,000 \times (1+0.1858) = \$1,126,510$.

11.2.4 Returns

Return on Debt

Table 7 presents the return on debt for our home sale. As with previous tables, each row corresponds to a different capital structure as indicated by the Leverage column. The next three columns correspond to the different sale prices - Low (\$0.8 million), Medium (\$1.1 million), and High (\$1.3 million). The numbers underneath these columns correspond to the **return on debt** measured in percent.

Leverage (%)	Sales Price			Expected	Risk	
	Low	Medium	High	Return	SD	Premium
0	N/A	N/A	N/A	N/A	N/A	N/A
25	5.00	5.00	5.00	5.00	0.00	0.00
50	5.00	5.00	5.00	5.00	0.00	0.00
75	5.00	5.00	5.00	5.00	0.00	0.00
95	-15.79	15.79	18.58	6.19	15.59	1.19

Table 7: One-Year Realized Returns, Expected Returns, Standard Deviations, and Risk Premium on Debt for Different Capital Structures. All Measures in Percent

With no debt, there is no return on debt, hence, the “N/A”s in the first row. Otherwise, the return to debtholders is computed as it always is - the money received divided by the money spent less one. The debtholder (i.e., bank) spends the principal amount of the loan,

²This calculation assumes that the loan is **non-recourse** - meaning the bank can only take the house if we don't pay them - all the proceeds from the sale go to the bank. Hence, the bank gets \$800,000 - the sales price - and the homeowner gets zero. In a recourse loan, the bank could go after other personal assets - bank accounts, cars, etc. - to recoup its investment.

which they give to the home buyer, and (hopefully) receives interest plus principal. Hence, the return on debt is

$$\text{Return on debt} = \frac{\text{Interest} + \text{Principal Repayment}}{\text{Principal}} - 1.$$

The fifth and sixth columns present the expected return and standard deviations (SD) for each set of realized returns. The last column presents the risk premium equal to the expected return minus the risk-free rate.

When leverage is low, the lender earns 5% on their investment *no matter what happens in the future*. That is, the lender is guaranteed a return of 5%. Consequently, the SD of the return on debt is zero, and the implied risk-free rate is 5%. The risk premia on debt for low leverage levels in this example are all zero because debtholders earn the same return no matter what happens in the future. In the real world, of course, we'd never get such a low rate on a loan because there is always some chance that we could default.

When leverage is 95%, the homeowner defaults on the loan in both the Low and Medium states as we saw above. In these two states, the lender only recovers the sale price of the home - \$800,000 or \$1,100,000. Therefore, the return on debt in the Low and Medium states are $(800,000 - 950,000)/950,000 = -15.79\%$ and $(1,100,000 - 950,000)/950,000 = 15.79\%$. In the high state, the lender is fully repaid and earns the interest rate on the loan, 18.58%.

The key takeaways from table 7 are as follows.

- When investors are guaranteed a return across all possible future states, that return is risk-free. Mathematically, the standard deviation of the return is zero.
- As leverage increases, the risk of default increases and, as long as the lender cannot fully recover what is owed, the expected return on the loan must increase to compensate lenders for the additional risk.³
- The expected return and the interest rate or yield on debt are not equal when there is a chance of default. The expected return will be less than the yield because the former is estimated with expected cash flows and the latter with promised cash flows. (Recall the discussion comparing risky bond yields and expected returns in chapter 7.)

Return on Equity

Table 8 presents the return on equity. The table structure is identical to that in Table 7. The homeowner receives the money from the home sale less what they owe on their loan - interest and principal.

$$\text{Return on equity} = \frac{\text{Sales Price} - \overbrace{(\text{Interest} + \text{Principal})}^{\text{\$ owed to lender}}}{\text{Equity investment}} - 1$$

³If the lender could fully recover what they are owed even in default, then default is economically meaningless and the expected return would be risk-free.

For example, if we sell the home for \$1.1 million with a \$750,000 mortgage, our return on equity is

$$\text{Return on equity} = \frac{1,100,000 - (37,500 + 750,000)}{250,000} - 1 = 25.00\%.$$

Leverage (%)	Sales Price			Expected Return	SD	Risk Premium
	Low	Medium	High			
0	-20.00	10.00	30.00	6.67	0.00	0.00
25	-28.33	11.67	38.33	7.22	27.40	2.22
50	-45.00	15.00	55.00	8.33	41.10	3.33
75	-95.00	25.00	105.00	11.67	82.19	6.67
95	-100.00	-100.00	246.96	15.65	163.56	10.65

Table 8: One-Year Realized Returns, Expected Returns, Standard Deviations, and Risk Premium on Equity for Different Capital Structures. All Measures in Percent

If we default on the loan - 95% leverage ratio and Low sales price - we lose all of our investment implying a return on equity of -100%, assuming the loan is non-recourse and the lender can't go after any other assets. This is the case for shareholders with **limited liability**. If the company in which we're a shareholder goes bankrupt, our losses are limited to our investment in the company. Thus, the worse return we could ever experience is -100%, which represents a total loss. However, if the lender had recourse against other assets or we as shareholders are not protected by limited liability, our return could be lower than -100%.

Table 8 shows that the return on equity is much more sensitive to capital structure than the return on debt. The standard deviation of returns is strictly increasing with leverage implying more debt means more risk for shareholders. Likewise, the expected return and risk premium are strictly increasing with leverage. These results imply two key takeaways.

- Shareholder risk, like creditor risk, is increasing in leverage - more debt means more volatile equity returns.
- Shareholder expected returns are increasing in leverage - more debt means higher expected equity returns.

Taking on more debt increases the risk to shareholders, who demand more compensation for the increased risk. Hence, shareholders' expected returns increase with leverage.

Return on Assets

Table 9 presents the return on assets when we sell the home. In this example, we have only one asset, the home, so the return on assets is the same as the return to operating assets.

$$\text{Return on assets} = \frac{\text{Sales Price}}{\text{Purchase Price}} - 1$$

For example, when the future sales price is Low, return on assets equals $800,000/1,000,000 - 1 = -20\%$. Likewise, when the future sales price is High, return on assets equals $1,300,000 / 1,000,000 - 1 = 30\%$.

Leverage (%)	Sales Price			Expected		Risk
	Low	Medium	High	Return	SD	Premium
0	-20.00	10.00	30.00	6.67	20.55	1.67
25	-20.00	10.00	30.00	6.67	20.55	1.67
50	-20.00	10.00	30.00	6.67	20.55	1.67
75	-20.00	10.00	30.00	6.67	20.55	1.67
95	-20.00	10.00	30.00	6.67	20.55	1.67

Table 9: One-Year Return on Assets (%)

The key takeaways from table 9 are as follows.

- The expected return on assets is unaffected by how we finance the home. Put differently, the value of assets is unaffected by capital structure.
- The risk, as measured by the standard deviation, of assets is also unaffected by how we finance the home.

Because capital structure is has no affect on the return on assets (or operating assets), we refer to the return on assets as the **unlevered cost of capital**. The unlevered cost of capital is the cost of capital if the asset was financed entirely with equity - no debt. Alternatively, the unlevered cost of capital is the cost of capital that is *unaffected* by leverage.

11.2.5 Summary

Let's summarize the lessons of this section.

1. Leverage does not affect the value or expected return of the house. In other words, O and r^O (or equivalently A and r^A because there is no excess cash) do not change as we change how we finance the asset. We can use equation 11.1 or equation 11.2 with $C = 0$ to verify this result.

$$\begin{aligned}
 \text{Leverage} = 0\%: & \quad r^A = 0.00 \times 0.05 + 1.00 \times 0.0667 = 0.0667 \\
 \text{Leverage} = 25\%: & \quad r^A = 0.25 \times 0.05 + 0.75 \times 0.0722 = 0.0667 \\
 \text{Leverage} = 50\%: & \quad r^A = 0.50 \times 0.05 + 0.50 \times 0.0833 = 0.0667 \\
 \text{Leverage} = 75\%: & \quad r^A = 0.75 \times 0.05 + 0.25 \times 0.1167 = 0.0667 \\
 \text{Leverage} = 95\%: & \quad r^A = 0.95 \times 0.0619 + 0.05 \times 0.1565 = 0.0667
 \end{aligned}$$

2. With no leverage, the expected return on equity equals the expected return on the asset - $r^E = r^A$. As leverage increases, the expected return to equity increases. We can use equation 11.4, or equation 11.5 with $C = 0$, to verify these results.

$$\begin{aligned} \text{Leverage} = 0\%: \quad r^E &= 0.0667 + \frac{0.00}{1.00} \times (0.0667 - 0.05) = 0.0667 \\ \text{Leverage} = 25\%: \quad r^E &= 0.0667 + \frac{0.25}{0.75} \times (0.0667 - 0.05) = 0.0722 \\ \text{Leverage} = 50\%: \quad r^E &= 0.0667 + \frac{0.50}{0.50} \times (0.0667 - 0.05) = 0.0833 \\ \text{Leverage} = 75\%: \quad r^E &= 0.0667 + \frac{0.75}{0.25} \times (0.0667 - 0.05) = 0.1167 \\ \text{Leverage} = 95\%: \quad r^E &= 0.0667 + \frac{0.95}{0.05} \times (0.0667 - 0.05) = 1.7000 \end{aligned}$$

3. As leverage increases, the expected return to debt, r^D , increases but more slowly than the expected return to equity. At 100% leverage, the expected return on debt equals the expected return on assets - $r^D = r^A$.

Figure 11.1 illustrates these lessons by plotting the return on assets, debt, and equity for different levels of leverage. While the exact numbers are specific to our home example, the broad patterns and relations between the lines hold for any assets and financial policy. The




Figure 11.1: Return on Assets, Debt, and Equity for Different Leverage Ratios, $D/(D+E)$

return on assets represented by the blue line is perfectly horizontal implying that the return

on assets is the same for all leverage ratios. The return on debt represented by the red line is flat for lower levels of leverage. At a leverage ratio of 76.2%, default risk kicks in and the cost of debt begins to increase. At 100% leverage, the firm is financed entirely with debt so $r^A = r^D$ and the red and blue line meet. With no debt financing - 0% leverage - the return on equity is equal to the return on assets, $r^E = r^A$. However, the return on equity increases quickly with leverage.

Looking at equation 11.1, it's natural to wonder: How can r^A remain constant when we change leverage, $D/(D + E)$, if r^A is a function of leverage? There are two offsetting effects when we increase leverage.

1. The return on equity (r^E) (and possibly debt, r^D) increases, which increases the return to assets (r^A).
2. The weight on equity ($E/(D + E)$) decreases, which decreases the return to assets (r^A).

These two effects cancel each other out resulting in a constant return on assets for all levels of leverage.

Another useful visual of these results is in Figure 11.2. The value and risk of a firm's assets are represented by the size of pie. How we finance those assets - debt and equity - corresponds to how we slice the pie. But, how we slice the pie doesn't change the size of the pie - value or risk. What slicing does change is who - debt or equity - receives the value and who bears the risk.

11.3 When is Financing Irrelevant for Value?

The lessons from the previous section follow from the work of Franco Modigliani and Merton Miller.⁴ Modigliani and Miller showed that under certain conditions the following results, known as **M&M's Propositions 1 and 2** or more simply **M&M**, would hold true.

1. The value and expected return of a firm's assets (A and r^A) are unaffected by the choice of how they are financed. As we vary the proportions of debt (D) and equity (E), the value and expected return on assets don't change because the underlying business risk is not changing.
2. The expected return on equity (r^E), and to a lesser degree the expected return on debt (r^D), increases as leverage ($D/(D + E)$) increases because the risk to equityholders, and debtholders, is increasing.

If we have excess cash on the balance sheet, then we can modify these two results to say the following.

⁴Franco Modigliani and Merton Miller's seminal 1958 paper, "The Cost of Capital, Corporation Finance and the Theory of Investment," American Economic Review 48, pages 261 - 297.




Figure 11.2: M&M Conservation of Value and Risk

1. The value and expected return of a firm's *operating assets* (O and r^O) are unaffected by the choice of how they are financed. As we vary the proportions of *net debt* ($D - C$) and equity (E), the value and expected return on *operating assets* don't change because the underlying business risk is not changing.
2. The expected return on equity (r^E), and to a lesser degree the expected return on debt (r^D), increases as *net leverage* ($(D - C)/Enterprise\ value$) increases because the risk to equityholders, and debtholders, is increasing.

A little thought reveals that these propositions have profound and wide-ranging implications. Because financial policy doesn't matter for value, CFOs and the financial sector more broadly are irrelevant. We might as well flip a coin when it comes to figuring out how we should finance a car, home, project, company, etc. But, this can't be true and, fortunately, isn't.

11.3.1 Perfect Capital Markets

The M&M propositions rely on a set of assumptions known as **perfect capital markets**. To clarify what each assumption means, we've included an example in the context of our home application.

1. No **transaction costs**. There are no costs (e.g., fees) to obtaining a mortgage.
2. No **taxes**. There is no effect of mortgage interest on the taxes we pay as homeowners or on the capital gains from the same of the home.
3. No **bankruptcy costs**. The bank and homeowner bear no costs if the homeowner defaults on the loan, things like legal fees or court costs.
4. No **information asymmetry**. The lender knows as much about the home as the homeowner, and vice versa.
5. No **agency conflicts**. The homeowner has the same interest in taking care of the home as the lender, regardless of the size of the mortgage.
6. No **capital market inefficiencies**. Loan interest rates accurately capture the risk of the borrower.

It should be clear from the examples that few if any of these assumptions are true, which begs the question: How in the world did this idea lead to a Nobel Prize?⁵ Answer: By knowing the conditions under which capital structure is *irrelevant* for firm value, we know the conditions under which it is *relevant*. For example, because mortgage interest is tax deductible, using debt can create value by reducing our taxes. Likewise, because bankruptcy is costly, using debt can destroy value by increasing the likelihood of bankruptcy. Thus, the insight of M&M can be summarized as follows.

If we want to use financial policy to create value for our company, we must be able to clearly articulate how that policy addresses one or more of the perfect capital markets assumptions.

Before exploring ways in which financial policy can create value, let's examine the implications of some financial policies in the context of perfect capital markets. Doing so allows us to introduce some new concepts, emphasize the mechanics of analyzing financial policies, and avoids falling prey to logical fallacies.

⁵Merton Miller was awarded the 1990 Nobel Prize in Economics for this work. His co-author, Franco Modigliani, had already won the award in 1985.

11.4 Leveraged Recaps, Payout Policy, and Excess Assets

What happens when a firm changes its capital structure? Let's answer this question using a hypothetical scenario in which NVIDIA Corporation (NVDA) - a producer and seller of computer chips - undertakes a **leveraged recapitalization**, or **leveraged recap**. A leveraged recap is when a firm issues debt and distributes the proceeds to shareholders either by repurchasing shares or paying a dividend. We want to understand what happens to NVIDIA when it engages in a \$500 billion leveraged recap when capital markets are perfect? More specifically, what happens to the value of the company? The share price? The cost of capital for debt and equity? The unlevered cost of capital?

11.4.1 The Perfect Capital Markets Benchmark

Let's start by answering these questions assuming that capital markets are perfect, i.e., the assumptions behind the M&M results hold. Doing so will establish a baseline and ensure we're performing calculations correctly because in an M&M world, capital structure and payout policy are both irrelevant for firm and shareholder value. All these policies do is allocate risk.

The left panel of Table 10 presents NVIDIA's market value balance sheet as of April 22, 2024 just before our hypothetical leveraged recap. From its accounting balance sheet, NVIDIA had \$11.1 billion in debt and leases, and \$25.9 billion in cash, cash equivalents, and short-term investments. Its market capitalization was \$1,905 billion. Because the MVBS has to balance, the value of operating assets - enterprise value - equals the value of equity plus debt minus cash or \$1,890.2 billion.

Underneath the MVBS is additional market information and cost of capitals. NVIDIA had 2.5 billion shares outstanding at a price per share of \$762. The company's leverage ratio, $11.1/1,961.1 = 0.58\%$ is only slightly above zero reflecting a predominantly equity-financed company. Its net leverage ratio is slightly negative, $(11.1-25.9)/1,890.2 = -0.79\%$, because it has more cash than debt. NVIDIA's debt cost of capital of 4.5% is estimated from information provided in its 2023 10-K filing. The 13.42% equity cost of capital comes from an application of the CAPM.⁶ The asset cost of capital, r^A , computed using equation 11.1 is

$$\begin{aligned} r^A &= \frac{Debt}{Debt + Equity} r^D + \frac{Equity}{Debt + Equity} r^E \\ &= \frac{11.1}{11.1 + 1,905.0} \times 4.50\% + \frac{1,905.0}{11.1 + 1,905.0} \times 13.42\% = 13.37\%. \end{aligned}$$

⁶As of April 22, 2024, the yield on a 30-year Treasury bond was 4.72% and NVIDIA's equity beta was 1.74 per Yahoo! Finance. Assuming a market risk premium of 5% produces an equity cost of capital of 13.42%.

Pre Debt Issue				Post Debt Issue			
Assets		Claims		Assets		Claims	
Cash (C)	25.9	Debt (D)	11.1	Cash (C)	525.9	Debt (D)	511.1
Operating (O)	1,890.2	Equity (E)	1,905.0	Operating (O)	1,890.2	Equity (E)	1,905.0
Total	1,916.1	Total	1,916.1	Total	2,416.1	Total	2,416.1
Price per share (\$)				762.00			
Shares outstanding (bil)				2.50			
Leverage ratio, D/D+E (%)				0.58			
Net leverage ratio, (D-C)/(E+D-C)				-0.79			
Cost of debt, r^D				4.50			
Cost of equity, r^E				13.42			
Cost of assets, r^A				13.37			
Cost of operating assets, r^O				13.49			
Price per share (\$)				762.00			
Shares outstanding (bil)				2.50			
Leverage ratio, D/D+E (%)				21.15			
Net leverage ratio, (D-C)/(E+D-C)				-0.79			
Cost of debt, r^D				6.95			
Cost of equity, r^E				13.44			
Cost of assets, r^A				12.07			
Cost of operating assets, r^O				13.49			

Table 10: NVIDIA's Market Value Balance Sheet as of April 22, 2024 Before and After Hypothetical Debt Issuance (\$bil)

NVIDIA's operating cost of capital is computed using equation 11.2.

$$\begin{aligned}
 r^O &= \frac{Debt - Cash}{Equity + Debt - Cash} r^D + \frac{Equity}{Equity + Debt - Cash} r^E \\
 &= \frac{11.1 - 25.9}{1,890.2} \times 0.045 + \frac{1,905.0}{1,890.2} \times 0.1342 = 13.49\%
 \end{aligned}$$

This operating cost of capital is higher than the asset cost of capital, which includes the moderating effect of cash.

What Happens Immediately After the Debt Issuance?

Immediately after the debt issuance, before NVIDIA distributes any money to shareholders, both debt and cash accounts increase by \$500 billion. As a result, **firm value**, which equals the sum of debt and equity, increases by the same amount. The firm is bigger, but the shareholders are no better off. Neither equity value nor enterprise value increased. Why? Capital markets are perfect so issuing debt is a zero NPV action. The firm gets \$500 billion in cash and promises over the term of the debt to make payments whose present value equals \$500 billion. Consequently, the share price is unchanged at \$762.

Also unchanged is the operating cost of capital, r^O . This is the essence of the M&M propositions. By changing the financing, we're not changing the business per se. So, r^O after the debt issuance remains at 13.49%. Likewise, because the firm has yet to distribute the cash, the increase in cash offsets the increase in debt dollar for dollar. Net leverage is unchanged at -0.79%.

What has changed is the firm's debt cost of capital, leverage ratio, and asset composition. The cost of the new debt is 7% because creditors know that the firm will be much riskier

after the debt issuance. Thus, the firm's total debt after the issuance is a portfolio of existing 5% debt and new 7% debt, and its total debt cost of capital is a value-weighted average of the two rates, 6.95%.⁷ From equation 11.5, we can see that the cost of equity has increased slightly.

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.10 + \frac{511.06 - 525.90}{1,905.00} (0.10 - 0.0695) = 13.44\%.$$

However, this increase only occurs because the cost of debt increased. Both the operating cost of capital and the net leverage ratio are unchanged as we have seen. In other words, the retention of cash offsets the large debt issuance.

Another way to understand the negligible effect of the debt issuance on equity risk is in the changing asset composition of the firm. Before the debt issuance, the firm's assets consisted of 1.35% cash and 98.65% operating. After the debt issuance, 21.77% of the firm's assets are cash and 78.23% are operating. Because of this shift to safe cash, the firm as a whole is less risky even though its operating risk hasn't changed - r^O is unchanged. As a result, the asset cost of capital declines from 13.37% to

$$r^A = \frac{525.9}{525.9 + 1,890.2} \times 6.95\% + \frac{1,890.2}{525.9 + 1,890.2} \times 13.49\% = 12.07\%.$$

Equity holders have a claim to a less risky firm even though their use of debt financing has increased. So, using equation 11.4, we get the same expected equity return as above.

$$r^E = r^A + \frac{Debt}{Equity} (r^A - r^D) = 0.1207 + \frac{511.06}{1,905.00} (0.1207 - 0.0695) = 13.44\%$$

The change in r^A may be confusing in light of the house example. However, in the house example, there was only one type of asset - the house. In other words, there was no excess cash, $C = 0$, and operating assets equaled total assets. What's different here is that we're not just changing leverage at this point, we're changing the composition of assets. Immediately after the debt issuance, the amount of excess cash on the balance sheet balloons. This compositional change is what affects the return on assets, not the change in leverage.

Leveraged Recaps and Payout Policy

Table 11 presents NVIDIA's market value balance sheets after distributing the cash to shareholders. The left panel shows what happens if the money is distributed by repurchasing shares; the right panel shows what happens if the money is distributed by paying a one-time special dividend. Let's discuss the stock buyback first then the dividend.

7

$$\frac{11.1}{11.1 + 500.00} \times 0.05 + \frac{500.00}{11.1 + 500.00} \times 0.07 = 6.95\%$$

Post Repurchase				Post Dividend			
Assets		Claims		Assets		Claims	
Cash (C)	25.9	Debt (D)	511.1	Cash (C)	25.9	Debt (D)	511.1
Operating (O)	1,890.2	Equity (E)	1,405.0	Operating (O)	1,890.2	Equity (E)	1,405.0
Total	1,916.1	Total	1,916.1	Total	1,916.1	Total	1,916.1
Price per share (\$)				762.00			
Shares outstanding (bil)				1.84			
Leverage ratio, D/D+E (%)				26.67			
Net leverage ratio, (D-C)/(E+D-C)				25.67			
Cost of debt, r^D				6.95			
Cost of equity, r^E				15.75			
Cost of assets, r^A				13.40			
Cost of operating assets, r^O				13.49			

Table 11: NVIDIA's Market Value Balance Sheet as of April 22, 2024 After Stock Buyback and Dividend (\$bil)

The firm repurchases \$500 billion in shares so cash decreases by the same amount. At a price of \$762 per share, the firm repurchases $500/762 = 0.656$ billion shares leaving 1.844 billion shares outstanding. Therefore, the market capitalization of NVIDIA after the repurchase is $762 \times 1.844 = \$1,405.0$ billion. The value of equity has fallen by the same amount as cash leaving the market value balance sheet in balance - assets equals claims. Are equity holders worse off? Not at all. Those shareholders that chose to sell or **tendered** their shares just exchanged their shares for cash at a fair price.

The decline in cash relative to operating assets means the total assets are now riskier and the expected return on assets increases from 12.08% to 13.40%.

$$r^A = \frac{\text{Cash}}{\text{Cash} + \text{Operating}} r^C + \frac{\text{Operating}}{\text{Cash} + \text{Operating}} r^O = \frac{25.9}{25.9 + 1,890.2} 0.0695 + \frac{1,890.2}{25.9 + 1,890.2} 0.1349 = 13.40\%$$

(Remember, we assume that the return on cash equals the return on debt and the return on operating assets is not affected by capital structure.) The reason why r^A after the distribution (13.40%) is different from r^A before the debt issuance (13.37%), even though the cash and operating asset values are identical, is that the return to excess cash is assumed to have changed from 4.50% to 7.00% to keep pace with the increased cost of debt. Yes, in reality this is unlikely. Nonetheless, it is a common assumption.

Both leverage and net leverage have increased significantly relative to their levels before the debt issuance. Therefore, the return on equity should increase. Equation 11.5 implies an equity cost of capital of 15.73%.

$$r^E = r^O + \frac{\text{Debt} - \text{Cash}}{\text{Equity}} (r^O - r^D) = 0.100 + \frac{511.1 - 25.9}{1,890.2} (0.100 - 7.00) = 15.73\%$$

Equation 11.4 implies the same.

$$r^E = r^A + \frac{Debt}{Equity} (r^A - r^D) = 13.40 + \frac{511.1}{1,405.0} (13.40 - 7.00) = 15.73\%$$

Once the company distributes the cash, the assets of the firm are riskier. Coupled with the increase in debt, shareholders expected returns increase.

If the company distributes the money by paying a dividend, the cash and equity values on the market value balance sheet both decline by \$500 billion as with the stock buyback. Because the market value balance sheet is the same whether the firm repurchases stock or pays a dividend, all of the cost of capitals - debt, equity, assets, and operating - are the same. What differs between the stock buyback and dividend is the stock price. The dividend leads to a reduction in share price, whereas the buyback did not affect the stock price.

The total value of equity decreases following a dividend, and the number of shares has stay the same. So, the price per share must decline. To understand this dynamic, recall the fundamental value relation. The value of equity is the sum of the present values of future cash flows to shareholders. Imagine we are nanoseconds from receiving our dividend, Div_0 . The value of a single share of stock would be the following.

$$Price_0 = Div_0 + \frac{Div_1}{(1+r^E)} + \frac{Div_2}{(1+r^E)^2} + \dots$$

As soon as we receive the dividend, Div_0 , the value of the share is

$$Price_0^* = \frac{Div_1}{(1+r^E)} + \frac{Div_2}{(1+r^E)^2} + \dots$$

The difference in the before and after dividend prices, $Price_0 - Price_0^*$ is Div_0 . In other words, the price decreases by Div_0 following payment of the dividend.⁸

Despite the price decline, shareholders are no worse (or better) off. Yes, the value of their shares have been reduced but only because the company gave them a bunch of cash in the form of a dividend. Whereas the stock buyback had shareholders exchange shares for cash, the dividend simply reduces the value of the shares by the amount of cash.

Perfect Capital Markets Summary

Let's summarize the lessons from this exercise that assumes *perfect capital markets*.

- There is no point to a leveraged recap when capital markets are perfect because the transaction has no impact on firm or shareholder value. It just reallocates risk. Likewise, there is no reason to fret over the decision between share buybacks and dividends. Both have the same impact on shareholders so payout policy is irrelevant just like capital structure when the assumptions of M&M are true.

⁸The stock price decline following dividend payments is known as the ex-day drop. See "Do Price Discreteness and Transactions Costs Affect Stock Returns? Comparing Ex-Dividend Pricing Before and After Decimalization" by John R. Gram, Roni Michael, and Michael R. Roberts in the 2003 **Journal of Finance** for more details.

- Increasing excess cash, all else equal, reduces the risk of the company's assets as a whole and r^A decreases.
- The operating assets, r^O , are unaffected by changes in excess cash or capital structure. This is what M&M tells us. How we slice up the claims between debt and equity is irrelevant for the business and therefore the risk of the business.
- Increasing net leverage will increase the return on equity, r^E , a.k.a., the **levered return**.
- Shareholders are indifferent between receiving a dividend or selling their shares back to the company. Both actions leave shareholders no better or worse off. They are zero NPV transactions.

11.4.2 Considering Market Imperfections

Now let's consider what happens in the real world where capital markets are imperfect. In other words, when we allow for taxes, bankruptcy costs, information asymmetry, etc., what happens to debt and equity values? What happens to debt, equity, and the unlevered cost of capitals? Quantitatively answering these questions is difficult because quantifying most market imperfections is difficult. So, the following sections quantify where possible and discuss the effects of different market imperfections on our leveraged recap example.

Transaction Costs

Issuing debt is costly. Consider our housing example and a mortgage. There is a meaningful opportunity cost associated with the homeowner's time spent gathering and preparing documents - bank statements, proof of employment and wages, loan documents, etc. Mortgages also often come with fees corresponding to a home appraisal, home inspection, loan origination, loan application, credit report, and documents and preparation. These costs make a mortgage less attractive than paying all cash for a home.

Corporations also face costs when raising money from outside investors - both debt and equity. There are fees paid to commercial and investment banks who manage and underwrite the fund raising process. Fees are also paid to lawyers, accountants, and sometimes exchanges. A study by Lee, Lochhead, Ritter, and Zhao estimated that the costs for issuing debt averaged between 1.3% and 8.75% of the proceeds from the issuance.⁹ The cost for issuing equity either in an **initial public offering (IPO)** or a **seasoned equity offering (SEO)** ranged from 3.15% to 16.96%. Variation in both equity and debt issuance costs was due primarily to the size of the issuance. The larger the issuance, the lower the fraction paid in fees. For example, a \$5 million IPO faced fees averaging 16.96% or \$0.8 million, whereas \$500 million IPO faced fees averaging 5.72% or \$28.6 million.

⁹I Lee, S Lochhead, J. Ritter, and Q. Zhao, 1996, "The costs of raising capital," *Journal of Financial Research* 19(1), 59-74.

Transaction costs alone create a clear preference ranking of financing sources for firms. At the top are internal funds because there are no transaction costs associated using cash. Next is debt followed by equity, which is the more expensive of the two external financing options. Interestingly, a similar **pecking order** can arise when investors are less well-informed than managers. That is, there is **information asymmetry** between the firm's managers and investors.¹⁰

We can account for these transaction costs in our analysis by deducting them from the shareholders who have to pay these costs. In NVIDIA's leveraged recap example, imagine that the cost to issue the debt was 0.5%. NVIDIA would have to pay $0.005 \times 500 = \$2.5$ billion in fees to the banks underwriting the issuance. So, the post-debt issue market value balance sheet would show a Cash account of \$523.4 billion and an Equity account of \$1,902.5 billion. If NVIDIA also faced a transaction cost to repurchase its equity, such as a bid-ask spread, the post repurchase Cash and Equity accounts would again be lower than what is reported Table 11 by the total amount of the cost.

Taxes

Interest expense generated by a corporation is tax deductible meaning corporations can reduce the taxes they pay by issuing debt and paying interest. Lower taxes means more money for investors - creditors and shareholders - and therefore more value. In effect, issuing debt reduces the government's stake in the company, assuming they tax corporations as they do in most countries.¹¹ Consider NVIDIA and the three different scenarios depicted in Table 12. NVIDIA's projected earnings before interest and taxes (EBIT) are \$45.37 billion and assumed unaffected by the amount of debt they carry.

	No Net Debt	Current Net Debt	Post-Recap Net Debt
EBIT	45.37	45.37	45.37
Net interest expense	0.00	(0.67)	33.70
Pre-tax income	45.37	46.03	11.67
Taxes	9.53	9.67	2.45
Earnings	35.84	36.37	9.22
Income to all investors	35.84	35.70	42.92
Interest tax shield	0.00	(0.14)	7.08

Table 12: NVIDIA's Interest Tax Shield at a 21% Tax Rate Under Different Capital Structures (\$bil)

¹⁰See Myers and Majluf, 1984, "Corporate financing and investment decisions when firms have information that investors do not have," *Journal of Financial Economics* 13, 187-221.

¹¹Examples of countries with no corporate tax as of 2024 include the Cayman Islands, Bermuda, Bahamas, Jersey, and the United Arab Emirates

The leftmost column presents NVIDIA's P&L if the company had not debt or excess cash, i.e., net debt equal to zero. In this scenario, NVIDIA pays 21% of its EBIT in taxes or $45.37 \times 0.12 = \$9.53$ billion. Investors, who consist of only shareholders, receive \$35.84 billion. The middle column presents NVIDIA's P&L under its current capital structure containing \$25.9 billion in cash and \$11.1 billion in debt. Because cash and debt are both assumed to earn the same 4.5% return, the firm's net interest expense is negative, i.e., it makes more in interest income than it pays in expense. This results in a higher pre-tax income and higher taxes but also higher earnings for shareholders. However, the total income available to *all* investors - equity and debt - is \$35.70 billion, less than if the firm had no cash or debt. Finally, the rightmost column shows that after the leveraged recap, earnings plummet to \$9.22 billion, but the income available to all investors has increased significantly to \$42.92 billion.

There are several takeaways from this analysis. First, differences in the income to all investors row arise from the interest tax shield provided by debt. The tax shield can be computed as the difference in income to all investors relative to the No Net Debt case or as the product of the net interest expense and the tax rate. In the middle column, the interest tax shield is negative because of the negative net debt, $11.1 - 25.9 = -\$14.8$. Currently, NVIDIA is operating in a tax inefficient manner. The government is taking a large share of the firm's value through taxes and we see that investors - debt and equity - are paying for it with less cash flow. After the recap, the government gets significantly less in taxes because NVIDIA is able to expense its new, massive interest bill. But, less for the government means more for investors, all else equal. So, from a firm level perspective, investors are better off paying less in taxes and financing with debt helps accomplish this objective.

Second, earnings are declining with increases in debt usage yet shareholders are, in fact, better off. Looking only at earnings to assess shareholder welfare makes two mistakes. First, the value of the firm is determined by the sum of debt and equity values. While earnings are going down, the amount of money available to *all* investors, interest plus earnings, is going up. In other words, the value of the firm is increasing. Second, only looking at earnings differences across different capital structures ignores the investment differences. With no debt, shareholders put up all of the money needed to generate operating earnings. Hence, they receive all of the after-tax operating earnings. With debt financing, shareholders only put up some of the money. Hence, they are entitled to only some of the after-tax operating earnings. By reducing taxes, debt increases what is available to shareholders. This intuition is illustrated in figure ?? where the use of debt financing reduces the share of the corporate pie available to the government and therefore increases the value to creditors and shareholders.

Third, because financial policy affects value when there are taxes, it must affect either cash flows or discount rates in our fundamental value relation, and it shouldn't matter which. That is, we should get the same answer regardless of which approach we take. Let's consider the discount rate approach here. We'll examine the cash flow approach later in this chapter. We can replace the cost of debt, r^D , in our pre-tax WACC formula (equation 11.1 or 11.2) with the after-tax cost of debt, $r^D(1 - TaxRate)$. In other words, the after-tax WACC for




Figure 11.3: Debt's Impact on Taxes

the firm (i.e., all assets) is

$$r^{WACC} = \frac{Debt}{Equity + Debt} r^D (1 - TaxRate) + \frac{Equity}{Equity + Debt} r^E. \quad (11.10)$$

Using equation 11.4, we can also express the after-tax WACC for the firm in terms of the return on assets and debt.

$$r^{WACC} = r^A - \frac{Debt}{Equity + Debt} \times TaxRate \times r^D \quad (11.11)$$

We use equation 11.10 or 11.11 to estimate the discount rate for the free cash flows generated by the firm or project, inclusive of any cash flows generated by excess assets.

The after-tax WACC for the enterprise (i.e., operating assets), assuming cash and debt earn the same expected return, is

$$r^{WACC} = \frac{Debt - Cash}{Equity + Debt - Cash} r^D (1 - TaxRate) + \frac{Equity}{Equity + Debt - Cash} r^E. \quad (11.12)$$

Using equation 11.5, we can also express the after-tax WACC for the enterprise in terms of the return on operating assets and debt.

$$r^{WACC} = r^O - \frac{Debt - Cash}{Equity + Debt - Cash} \times TaxRate \times r^D \quad (11.13)$$

We use equation 11.12 or 11.13 to estimate the discount rate for the free cash flows generated by the enterprise or operating assets, ignoring any cash flows generated by excess assets.

In practice, people drop the “after-tax” descriptor and just refer to equations ?? through ?? as “the WACC.” This is sloppy and incorrect. There are many weighted average costs of capitals, yet the language persists. The key point is that these equations are used to discount future cash flows to compute value *inclusive of the interest tax shield*. Put differently, the

WACC is less than either the asset or operating cost of capital when there is leverage because of the interest tax shield. This lower cost of capital results in a higher valuation.

Prior to the leveraged recap, NVIDIA's post-tax WACC is

$$r^{WACC} = \frac{11.1 - 25.9}{1,890.2} \times 0.045 \times (1 - 0.21) + \frac{1,905.0}{1,890.2} \times 0.1342 = 13.50\%.$$

After the leveraged recap, NVIDIA's post-tax WACC is

$$r^{WACC} = \frac{511.1 - 25.9}{1,890.2} \times 0.045 \times (1 - 0.21) + \frac{1,405.0}{1,890.2} \times 0.1342 = 13.12\%.$$

The leveraged recap creates value by reducing the cost of capital, which is due to the reduction of taxes arising from interest expense.




Figure 11.4: Returns for Different Leverage Ratios, $D/(D+E)$

Figure 11.4 revisits our house by adding two lines to Figure 11.1. The dashed purple line is the after-tax cost of debt at a 36% personal income tax rate. Thus, the after tax cost of debt at relatively low levels of leverage is $5 \times (1 - 0.36) = 3.20\%$. This number increases at higher levels as the debt becomes riskier but remains below the pre-tax cost of debt - the red line. The dashed turquoise line corresponds to the after-tax WACC. Notice how the after-tax WACC decreases with leverage, in contrast to the pre-tax WACC or unlevered cost of capital which remains constant as leverage increases (the solid blue line). Thus, as leverage increases, the cost of capital decreases because of the interest tax shield, and value increases.

Figure 11.4 suggests that firms should be 100% levered - financed entirely with debt - because that is where the WACC is lowest. There are several practical limitations to this implication beginning with the amount of operating income. What companies should really do is issue enough debt to eliminate all taxes, which occurs when interest expense is equal to EBIT. In our NVIDIA example, we want

$$\underbrace{r_{Existing}^D Debt^{Existing} + r_{New}^D Debt^{New}}_{\text{Total interest expense}} = EBIT,$$

where we've distinguished between any existing debt ($D^{Existing}$) and its cost ($r_{Existing}^D$) from any new debt (D^{New}) and its cost (r_{New}^D). Solving for the amount of new debt using NVIDIA's values yields

$$\begin{aligned} Debt^{New} &= \frac{EBIT - r_{Existing}^D Debt^{Existing}}{r_{New}^D} \\ &= \frac{45.37 - 0.0695 \times 11.1}{0.070} \\ &= \$637.12. \end{aligned}$$

NVIDIA should issue an additional \$637.12 billion in debt at 7% to completely exhaust their projected taxable income.

Why doesn't NVIDIA do this? That is, why does NVIDIA pay any taxes when it can simply issue debt to eliminate them? Figure 11.5 shows that NVIDIA is not alone. The figure presents annual quartiles - 25th, 50th, and 75th percentiles - of the interest expense-to-EBIT ratio for U.S. publicly traded firms from 1950 to 2022. If firms were issuing enough debt for the interest expense to equal EBIT, this ratio would be 1.0. However, for the large majority of firms, this ratio is well below 1.0. In fact, for 75% of firms, the ratio never reaches 0.5 in any year during this 72-year window.

There are several reasons why firms don't issue debt so that their interest expense equals EBIT. First, there is uncertainty around next year's EBIT. So, firms can only choose their debt today based on an estimate of what EBIT will be. If their estimate is wrong, as it will almost surely be, they may have insufficient or excess interest. Second, since 2022, the tax law in the U.S. limits the amount of interest that firms with more than \$25 million in revenue can deduct to 30% of operating earnings. Before the leveraged recap, this limit wasn't binding. NVIDIA had \$0.50 billion of interest, which is 1.1% of EBIT. After the leveraged recap, interest is \$35.50 billion or 78.25% of EBIT. So, NVIDIA would only be able to deduct $0.30 \times 45.37 = \$13.61$ billion of interest from EBIT to compute the income base against which taxes are computed. That is, taxes would be computed as 21% of \$31.76 billion ($45.37 - 13.61$), which equals \$6.67 billion. In this case, NVIDIA has $35.51 - 13.61 = \$21.89$ billion of *excess interest* from a tax standpoint. Fortunately, it can use this excess interest expense to offset future taxable income, again subject to the total interest expense being less than 30% of EBIT. Third, personal taxes can mitigate the tax benefit of debt. Though U.S. taxes have varied greatly historically, typically ordinary income has been taxed at a higher rate than dividends and capital gains. In other words, while the company prefers




Figure 11.5: Returns for Different Leverage Ratios, $D/(D+E)$

to pay interest because it reduces corporate taxes, investors prefer to receive dividends and capital gains because of lower personal taxes. See the technical appendix for details. Finally, there are costs associated with issuing debt

Bankruptcy Costs

When firms issue debt, they risk default or bankruptcy. The more debt they use the more likely they are to default. We saw this in the housing example. With a 95% mortgage, there were some states of the world in which we were unable to repay the mortgage because the value of our home that we sold was less than what we owed to the bank. So, we defaulted on the mortgage and, in most cases, would have to declare bankruptcy. However, in the housing example, default wasn't a big deal because there were no costs associated with default. When the homeowner defaulted, we assumed the homeowner simply walked away, and the lender seized the home and received the fair market value. All the lender had to do was charge a higher interest rate to compensate for the additional risk so they could earn a fair rate of return in expectation (i.e., on average).

But, imagine that if we defaulted on our mortgage, we faced a bunch of additional costs. For example, the lender may sue us in which case we would have to pay legal fees, court costs, and potentially damages. Additionally, our credit score would decline significantly, which would hinder our access to and increase the cost of credit. In other words, bankruptcy in the real world is very costly for homeowners. Likewise, before defaulting, many homeowners do not have much incentive to take care of the house because they know the bank will take

it. When the bank does take the house, they face a number of fees and legal expenses. They also have to sell the house and face the transaction costs of doing so (e.g., realtor fees). The point is that there are many costs associated with bankruptcy that temper the use of debt by homeowners and the supply of debt by lenders. Bankruptcy is costly and too much debt increases the likelihood of experiencing those costs.

The logic is similar for corporations. Too much debt financing risks default and bankruptcy which is costly. There are two types of bankruptcy costs firms face. There are **direct costs of bankruptcy** that include court costs; fees for lawyers, accountants, consultants, and bankers; administrative costs; and liquidation or reorganization costs. Direct costs have been estimated to be 3% of the firm's market value of the company.¹² There are **indirect costs of bankruptcy** that occur prior to the actual declaration of bankruptcy. When firms become **financially distressed** or approaching bankruptcy, they experience costs often in the form of strained business relationships. For example, as a customer, we wouldn't want to buy a car from a manufacturer that was likely to go out of business. What if there is a problem with the car? Who would service it? From where would they get parts or trained technicians? Similarly, as a supplier, we probably wouldn't extend credit to a car manufacturer about to go out of business either. We may require upfront, cash payments which can be quite costly for the manufacturer in terms of pressure on their net working capital and free cash flow. Indirect costs have been estimated to be between 10% and 20% of the firm's market value. Figure 11.6 illustrates this intuition by showing how using debt can reduce the government's slice of the pie by reducing taxes while creating (or increasing) the slice of the pie going to the beneficiaries of default (lawyers, bankers, accountants, customers, suppliers).

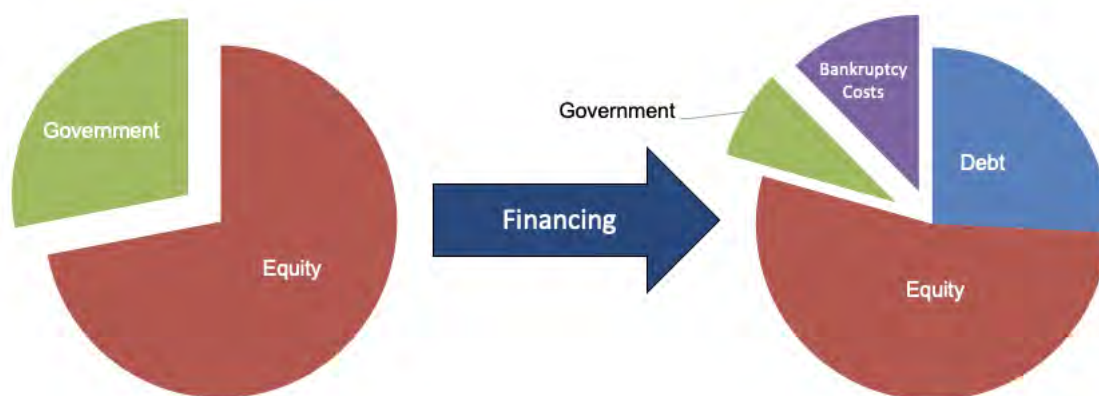


Figure 11.6: Debt's Impact on Reducing Taxes and Increasing Bankruptcy Costs

If we consider both taxes and bankruptcy costs in our capital structure decision, we arrive at what academics call the Tradeoff Theory of Capital Structure. This name is a little odd since virtually every theory in economics is a tradeoff of costs and benefits. Nonetheless, the basic idea of the Tradeoff Theory is that there is an optimal level of leverage that balances

¹²See for example Edward Altman, 1984, A Further Empirical Investigation of the Bankruptcy Cost Question, *Journal of Finance*, 1067-1089.

the tax benefits with the expected bankruptcy costs. This idea is illustrated in Figure 11.7. As we move from left to right along the horizontal axis, leverage increases. As leverage increases, the cost of capital (red curve) decreases and value increases (blue curve) because of the interest tax shield. The rates of decrease in the cost of capital and increase in value slows as leverage increases because expected bankruptcy costs grow and begin to offset the tax shield. At the point L^* , the cost of capital is at its minimum and value at its maximum. This the optimal capital structure for our imaginary firm. Leverage above L^* is suboptimal because value is decreasing and the cost of capital increasing. To the right of L^* , expected bankruptcy costs are greater than the tax benefit from debt.

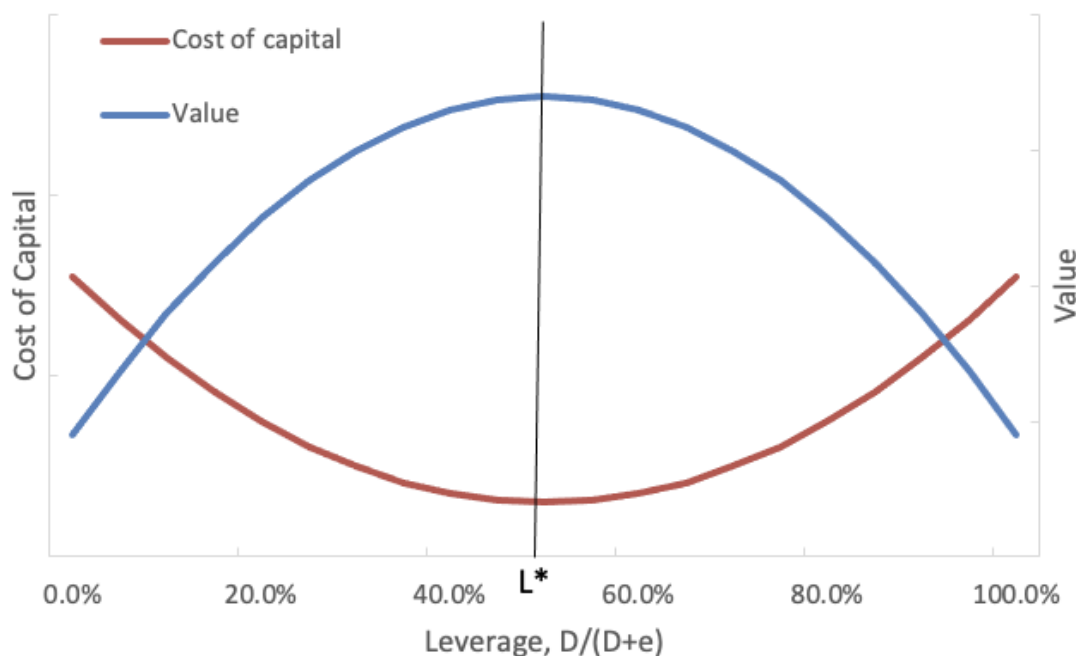


Figure 11.7: Optimal Capital Structure

The question is: How can we account for bankruptcy costs? Relative to transaction costs and taxes, measuring bankruptcy costs is significantly more difficult. Academic estimates are useful guides but statistically noisy. That is, the estimates are accompanied by wide confidence intervals because expected default costs vary across firms and time. Suggestive evidence of this variation is presented in Figures 11.8 and 11.9.

Figure 11.8 shows that leverage, as measured by the ratio of total debt to the sum of total debt and market capitalization, varies dramatically across industries. If we think about what differentiates high and low levered industries, an important aspect is **collateral**. Air transport, real estate, furniture and fixtures are business with significant tangible assets, assets that can be sold by creditors in bankruptcy. This is in contrast to the several service industries appearing towards the bottom of the graph. The primary assets of these businesses are people who will simply leave the company should it fall into bankruptcy. Of course, collateral differences aren't the only differences across these industries or the only reason by



Figure 11.8: Industry Median Leverage Ratios for 2022

they have different capital structures. But, it does represent an important difference because collateral helps limit lenders' losses in bankruptcy.

Figure 11.8 focuses on the health services industry and shows that over time an indus-

try's capital structure can vary significantly with economic conditions. In 2015, the median leverage ratio was 31.3% but only 14.6% three years later - less than half. In addition to increasing equity values, this decline in leverage reflects several changes in the industry including an increase in new firms that tend to be smaller and less levered.




Figure 11.9: Health Services Industry Median Leverage Ratio

Agency Problems

Another challenge with identifying an optimal capital structure is the incentive conflict between managers, creditors, and shareholders. All three groups have different incentives and capital structure can either exacerbate or mitigate the problems that arise from these different incentives.

Take managers who tend to be less risk tolerant than investors because their labor income and human capital is tied up in the firm. If the firm goes bankrupt, managers lose their jobs, income, and possibly wealth if they receive stock-based compensation (e.g., options and stock grants). Managers may also get personal benefits from their job. Perhaps they like to travel on private jets or buy expensive furniture for their offices. Maybe they like the power of running a large company and engage in a lot of questionable acquisitions. The problem is that managers may take actions that benefits themselves at the expense of investors, primarily shareholders.

Likewise, creditors are relatively risk averse as well. They care only about getting what they're promised because they don't reap the rewards if the firm performs exceptionally well. They do, however, eat all the losses should the firm struggle. The most creditors can hope for

is that they are paid what they've been promised. This is in sharp contrast to shareholders who get all the upside but suffer a limited downside. The most shareholders can lose is their entire investment, admittedly painful, but the most they can make is unlimited. The difference in payoffs between the two claimants is illustrated in Figure 11.10.




Figure 11.10: Value of Equity and Debt as a Function of Asset Value

If the value of assets is below the amount of debt, 100, equity is worthless and creditors get whatever is leftover. Once assets increase past 100, equity gets all the upside, which is capped for debt. In other words, as long as the company can repay its debt, creditors don't benefit from any additional performance which all goes to shareholders.

So, a major concern among creditors is that shareholders undertake actions that benefit them at creditors' expense. For example, shareholders could pay themselves excessive dividends. They could sell assets - collateral backing debt - to keep the firm afloat. They could take on excessively risky projects. They could issue additional debt that is senior in priority to the existing debt on the balance sheet. There are countless ways in which shareholders can hurt the value of creditors' claims. As a result, debt contracts are filled with provisions or **covenants** that restrict what borrowers can do (negative covenants) and detail what they must do (affirmative covenants). Violation of a covenant leads to a **technical default** in which case the lender often has the right to accelerate the debt due and demand immediate repayment, as well as terminating any future borrowing under a credit line or delayed draw term loan, for example.

The use of debt has several effects. Debt can exacerbate the tension between managers and shareholders by making managers even less willing to take risks. Debt can also discipline managers and mitigate **free cash flow problems** in which earnings are spent wastefully

on negative NPV projects. Likewise, too much debt can exacerbate the incentive conflict between creditors and shareholders. If firms are close to financial distress, shareholders have a strong incentive to gamble with creditors money by undertaking very risky projects because shareholders get all the upside and face limited downside. Too much debt can also lead to **underinvestment** if too much of the benefits go to paying off debt as opposed to shareholders.

Like bankruptcy costs, it's difficult to quantify agency costs. At this point, the best we can do is be aware of the costs and mitigate them through a judicious choice of debt and hiring smart lawyers to write and review the debt contracts.

11.4.3 Summary

The key points of our leveraged recap with market imperfections are the following.

- Transaction costs associated with the raising capital reduce the value of equity and create a preference ranking over financing sources: internal funds first, debt second, and equity last.
- Taxes reduce the cost of debt capital relative to equity capital because interest is tax-deductible. The after-tax WACC impounds the interest tax shield into the discount rate.

$$\begin{aligned} \text{Firm WACC} &= \frac{\text{Debt}}{\text{Debt} + \text{Equity}} r^D (1 - \text{TaxRate}) + \frac{\text{Equity}}{\text{Equity} + \text{Debt}} r^E \\ &= r^A - \frac{\text{Debt}}{\text{Equity} + \text{Debt}} (1 - \text{TaxRate}) r^D \\ \text{Enterprise WACC} &= \frac{\text{Debt} - \text{Cash}}{\text{Equity} + \text{Debt} - \text{Cash}} r^D (1 - \text{TaxRate}) + \frac{\text{Equity}}{\text{Equity} + \text{Debt} - \text{Cash}} r^E \\ &= r^A - \frac{\text{Debt} - \text{Cash}}{\text{Equity} + \text{Debt} - \text{Cash}} (1 - \text{TaxRate}) r^D \end{aligned}$$

We use the firm WACC to discount free cash flows generated by *all* of the firm's assets - excess and operating. We use the enterprise WACC to discount the free cash flows generated by the *operating* assets of the firm. Because these WACCs change with changes in (net) leverage, we must either assume that firms maintain a constant (net) leverage ratio or be prepared to recompute the WACC each time it changes.

- While taxes incentive the use of debt financing, bankruptcy costs dissuade the use of financing. Direct costs such as legal fees, court costs, etc. and indirect costs such as customer and supplier flight can be substantial. Collateral, or assets that can be used to secure the debt, can mitigate these costs and reduce the debt cost of capital. Too much debt can also hamper growth when too much of a firm's operating earnings are being used to payback creditors as opposed to reinvesting in profitable investment opportunities.
- Optimal capital structure balances the costs and benefits of debt financing to reduce the cost of capital to its lowest point and maximize the value of the firm or project.

11.5 Project Financing's Impact on the Firm

Now let's explore what happens to a firm when it takes on a project and must choose among different financial strategies. As with the leveraged recap, we'll start in a perfect capital markets setting to establish a baseline before incorporating market imperfections. We'll also keep our firm and project simple to avoid distractions that have no impact on the lessons of this example.

11.5.1 Firm and Project Setup

	Financial Strategies					
	Pre-Issue	All Equity	All Debt	All Cash	50% Cash 50% Debt	45% Debt 55% Equity
Assets - Cash						
Cash	20.00	20.00	20.00	10.00	15.00	20.00
Expected return on cash, r^C	5.00	5.00	5.00	5.00	5.00	5.00
Interest income, $r^C \times C$	1.00	1.00	1.00	0.50	0.75	1.00
Assets - Operating						
EBIT	10.00	11.50	11.50	11.50	11.50	11.50
Operating cost of capital, r^O	10.00	10.00	10.00	10.00	10.00	10.00
Enterprise value, $O = EBIT/r^O$	100.00	115.00	115.00	115.00	115.00	115.00
Claims - Debt						
Debt value, D	50.00	50.00	60.00	50.00	55.00	54.50
Debt cost of capital, r^D (%)	5.00	5.00	5.00	5.00	5.00	5.00
Interest, $r^D \times D$	2.50	2.50	3.00	2.50	2.75	2.73
Claims - Equity						
Equity value, $E = V - D$	70.00	85.00	75.00	75.00	75.00	80.50
Equity cost of capital, r^E (%)	12.14	11.76	12.67	12.67	12.67	12.14
Earnings, Earn = EBIT - Interest	8.50	10.00	9.50	9.50	9.50	9.78
Shares outstanding (mil)	5.00	5.67	5.00	5.00	5.00	5.37
Price/share, P	14.00	15.00	15.00	15.00	15.00	15.00
Earnings per share, EPS=Earn/Shares	1.70	1.76	1.90	1.90	1.90	1.82
Price-to-earnings ratio, P/E=P/Earn	8.24	8.50	7.89	7.89	7.89	8.24
Firm						
Firm value, $V = C + O$	120.00	135.00	135.00	125.00	130.00	135.00
Asset cost of capital, r^A	9.17	9.26	9.26	9.60	9.42	9.26

Table 13: Pre- and Post-Project-Financing Information in Perfect Capital Markets. All Dollar Values are in Millions Except Per Share Figures

The leftmost column - Pre-Issue - of Table 13 presents the market value balance sheet and some additional information for our firm prior to the announcement of the project and any financing. Our firm has \$20 million in cash earning 5% while its operating assets generate \$10 million per year in EBIT. Assuming this firm has no working capital or long-term investment needs, then EBIT is the equal to the unlevered free cash flow. (Remember, there are no taxes in perfect capital markets.) With an operating cost of capital equal to 10%, the enterprise value of our firm can be computed as the present value of a perpetuity, $10/0.10 = \$100$. The sum of cash and operating assets equal firm value, or \$120 million.

Our firm has financed these assets with \$50 million in debt at a 5% cost of capital. Deducting the debt from the value of the firm implies $120-50=\$70$ million of equity value or market capitalization. With five million shares outstanding, the price per share is $70/5 = \$14$ per share. Earnings equal EBIT less *net* interest expense or $10 - (2.50 - 1.00) = \$8.50$ million. Thus, earnings per share (EPS) are $8.50 / 5.00 = \$1.70$, and the price-to-earnings (P/E) ratio is $14.00/1.70 = 8.24$.

The equity cost of capital can be computed several ways beginning with equations 11.4 and 11.5.

$$r^E = r^A + \frac{Debt}{Equity} (r^A - r^D) = 0.0917 + \frac{50}{70} (0.0917 - 0.05) = 12.14\%$$

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.10 + \frac{50 - 20}{70} (0.10 - 0.05) = 12.14\%$$

The asset expected return, r^A is the value weighted sum of cash and operating asset expected returns.

$$r^A = \frac{Cash}{Cash + Operating} r^C = \frac{Operating}{Cash + Operating} r^O = \frac{20}{20 + 100} \times 0.05 + \frac{100}{20 + 100} \times 0.10 = 9.17\%$$

The project is expected to generate \$1.50 million in unlevered free cash flow in perpetuity and has a risk similar to that of the operating assets of the firm. Discounting these cash flows by the operating cost of capital, r^O means the project value is $1.50 / 0.10 = \$15$ million. Deducting the upfront cost of \$10 million implies a project NPV of \$5.00 million.

Now, let's see what happens to the firm under different financing strategies for the project. As with

11.5.2 The Perfect Capital Markets Benchmark

Equity Financing

The second column labeled "All Equity" shows what happens when the firm raises \$10 million of equity to fund the project by issuing 0.67 million shares at a price of \$15 per. (We'll see why \$15 per share in a moment.) Both enterprise and firm value increases by

the value of the project, \$15 million. This is confirmed by discounting the new EBIT of $10 + 1.50 = 11.50$ by the operating cost of capital, $11.50 / 0.10 = \$115$ to obtain the enterprise value. Add in the cash, \$20 million, to get the firm value of \$135 million. Because operating assets increased and cash did not, the asset cost of capital increased from 9.17% to 9.26%. In effect, the firm tilted its asset portfolio away from safe cash and towards riskier operating assets.

Nothing concerning the firm's cash or debt changes; the project is financed entirely with new equity. Equity value increases by the entire value of the project because shareholders paid for the entire project. This can be seen by subtracting the debt value from the firm value, $135 - 50 = \$85$ million. It can also be seen by discounting the new earnings stream, \$10 million, by the equity cost of capital. However, we can't use the old equity cost of capital, 12.14%, because both the financing and project NPV have changed the capital structure and therefore the equity cost of capital. The new equity cost of capital is

$$r^E = r^O - \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.10 + \frac{50 - 20}{85} (0.10 - 0.05) = 11.76\%.$$

The equity issuance and the project NPV increase the amount of equity relative to debt and therefore reduce the firm's leverage. Less leverage means less risk for shareholders, all else equal. Discounting the new earnings stream with this new equity cost of capital confirms that the equity value is, $\$10.00 / 0.1176 = \85 million.

Finally, shares outstanding increases by the newly issued 0.67 million shares implying that the new price per share is $\$85 \text{ million} / 5.67 \text{ million} = \15.00 . Earnings per share is \$1.76, slightly higher than before the stock issuance, and the price-to-earnings (P/E) ratio has increased from 8.24 to 8.50. Notice that the even though the firm issued equity, earnings per share increased. In other words, there doesn't seem to be any **dilution** - reduction in EPS - as a result of the share issuance. Practitioners and financial journalist often assume that share issuances are bad - value-destructive - for shareholders because the new shares dilute those of the existing shareholders. In other words, the slice of the pie existing shareholders hold gets smaller with share issuances. This conclusion makes two mistakes. First, it assumes shareholders care more about earnings per share than price per share; they do not. Second, it ignores what's done with the money once it is raised. As long as the firm does not invest the money in negative NPV projects, no shareholders - new or old - will be harmed by an equity issuance. In fact, EPS can go up or down depending on the project NPV, but as long as that NPV is positive, shareholders are strictly better off.

Now let's understand why the new shares had to be issued at \$15 per share. Figure 11.11 illustrates how the project NPV is divided between existing (red line) and new (blue line) shareholders for different issuance prices. The value gained for each group is measured by comparing the value of their shares before and after the issuance. For example, at point C, new shareholders get none of the project value and existing shareholders get all of it - \$5 million NPV. Before the issuance, existing shareholder value was \$70 million. After the issuance, the value is $\$15.00 \times 5.00 = \75.00 million. Existing shareholders gain \$5.00 million, which is equal to the NPV of the project. New shareholders paid \$15.00 per share to purchase shares that were worth just that, \$15.00. So, new shareholders gain (and lose)

nothing from purchasing the shares. It's a zero NPV investment or a fair deal for new shareholders. At point A, new shareholders get all of the gains from the project, \$5 million, and existing shareholders get none. And, at Point B, new and existing shareholders split the project NPV equally - \$2.5 million each - at an issuance price of \$11.60.




Figure 11.11: Division of NPV Between Existing and New Shareholders for Different Issue Prices

Looking at figure 11.11, existing shareholders would never issue shares at a price lower than the \$9.33 corresponding to point A because they would lose money (negative value gain) from issuing shares too cheaply. Likewise, new shareholders would never pay more than the \$15.00 corresponding to point C because they would lose money (negative value gain) from buying overpriced shares. We might think that the ultimate price would fall somewhere in the range between \$9.33 (point A) and \$15.00 (point C) depending on the relative bargaining power of the two groups of shareholders. However, Figure 11.12 shows that as long as markets are competitive, the share issuance price has to be \$15.00. Figure 11.12 plots the issuance share price against the final share price. The black dashed 45 degree line identifies all points at which the issuance price and the final share price are equal.

Any issuance price less than \$15.00 means that new shareholders can buy the stock and immediately earn money when the stock price incorporates the value of the project. Take point A from the previous figure, for example. At this point, new investors buy shares of the firm for \$9.33, immediately after which the price increases to \$15.00. Knowing this price increase is coming, more investors will want to purchase the stock, driving up demand and the price. Investors will keep demanding the stock until they are indifferent between buying and not buying the stock, which occurs when the issue price equals the final price, which in this case is \$15 per share. Similarly, if new investors were offered shares at a price greater




Figure 11.12: Final Share Price for Different Issue Prices

than \$15.00, no one would buy the stock because immediately after doing so the share price would fall. So, with perfect capital markets, the only price at which the firm can issue shares is the price at which existing shareholders get all of the benefits from the new project.

Debt Financing

The third column of Table 13 labeled “All Debt” shows the consequences of financing the project entirely with debt costing 5%.¹³ Enterprise and firm value increase by the value of the project, \$ 15 million, while debt increases by the cost of the project, \$10 million. The increased debt leads to a larger interest expense, \$3.0 million, and consequently lower earnings than when the project was financed with equity. The value of equity is the difference between the value of the firm and its debt, or $135 - 60 = \$75$ million. Equivalently, the equity value equals the sum of discounted earnings at the new equity cost of capital.

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.10 + \frac{60 - 20}{75} (0.10 - 0.05) = 12.67\%$$

¹³If the new debt was issued at a different interest rate, the interest rate on all of the debt would be a value-weighted sum of interest rates. For example, if the new debt was issued at a 7% interest rate, then the interest rate on all of the firm’s debt would be

$$r^D = \frac{50}{60} 0.05 + \frac{10}{60} 0.07 = 5.33\%.$$

That is, $9.50 / 0.1267 = \$75$ million. The impact of the debt issuance on the equity cost of capital is muted by the increase in the value of the equity from the project NPV. While the value of debt increased by \$10 million to fund the project, the value of equity increased by \$5 million because the project is positive NPV. In other words, leverage increased by less than it would have had the project been zero NPV. Consequently, the cost of equity did not increase by as much.

Earnings per share increases from \$1.70 to \$1.90 cents because the project increased earnings without increasing the number of shares outstanding. This might lead one to believe that debt financing is preferred to equity financing because of the larger impact on EPS compared to the equity issuance. This is a mistake. The higher earnings per share in this instance coincides with an increase in the risk of those earnings caused by the increase in leverage. In the context of our fundamental value relation, each shareholder expects to receive higher cash flows - numerators are bigger - but discounted at a higher rate - denominator is bigger. Indeed, $1.90 / 0.1267 = \$15.00$ per share when we finance with debt and $1.76 / 0.1176 = \$15.00$ per share when we finance with equity. The stock price is the same despite the different EPS.

Cash Financing

The fourth column labeled "All Cash" shows that the impact on shareholders is identical to using all debt. Equity value, equity cost of capital, price per share, earnings, EPS, and P/E ratio are all the same with these two financing strategies. Why? Because reducing cash is equivalent to increasing debt when the two have the same expected return. The firm can issue a dollar of debt and pay interest or it can reduce a dollar of cash and earn less interest. The effect on net leverage and earnings is the same.

That said, instead of debt increasing and cash staying constant, the exact opposite occurs when financing with cash. As a result, the composition of assets is different when financing with debt and cash (and equity). Financing with cash reduces the cash balance, which when coupled with the increase in operating assets from the project, tilts the assets much towards risky operating assets and away from safe cash. As a result, the asset cost of capital increases to its highest level, 9.60%. This increase in asset risk is completely borne by shareholders just as if they had financed the firm with debt. Except instead of the risk coming from new debt, it comes from a riskier asset portfolio. Also, the size of the firm is smaller when financing with cash because shareholders are effectively selling off some of the firm's assets - cash - to invest. The firm spends \$10 million in cash to grow by \$15 million for a net effect of \$5 million in growth. When raising external capital - debt or equity - the firm grows by the \$10 million raised and the \$5 million NPV of the project for a net effect of \$15 million in growth. But, don't be fooled into thinking bigger is better! Shareholders are indifferent are no better or worse off across these different strategies. Look at the share price!

Mixed Financing

The last two columns of Table 13 examine the implications of using multiple sources of financing. The “50% Cash/50% debt” column presents a middle ground between the all cash and all debt strategies. The implications for shareholders is unchanged from all debt and all cash financing and the 50-50 mix is not special. We get the same implications for shareholders regardless of the proportions of cash and debt. As we discussed, shareholder KPIs are unaffected because cash and debt are two sides of the same coin. Issuing debt is equivalent to spending cash in terms of the impact on net leverage and earnings. What differs is the amount of cash and debt and, as a result, the size and risk of the firm as a whole.

The last column of Table 13 looks at financing the project with 45% debt and 55% equity. These proportions are special because they ensure the project is financed in the same manner as the firm. The project net leverage ratio is $4.50/15.00 = 30\%$, the same as the firm before the issuance. Because the project and firm financing are the same, the net leverage ratio of the firm after the issuance is unchanged.

$$\begin{aligned}\text{Pre-issue net leverage} &= \frac{50 - 20}{100} = 30\% \\ \text{Post-issue net leverage} &= \frac{54.50 - 20}{115} = 30\%\end{aligned}$$

Because the net leverage is the same, the equity cost of capital is the same, 12.14%, and the price to earnings ratio is the same, 8.24.

Maintaining a constant capital structure makes capital budgeting easier because the cost of capitals don't change. In other words, the discount rate we use for the firm and project are the same. Additionally, a constant capital structure is loosely consistent with empirical evidence. Firms tend to maintain a loose target when it comes to their capital structure. Large swings in capital structure are relatively infrequent and often temporary.¹⁴

Note, we don't have to finance the project with the exact proportions of debt and equity as the firm to maintain a constant capital structure. The leverage that matters for estimating the project cost of capital is the *incremental* leverage. So, the firm can finance the project one way and then take other steps to bring its capital structure back in line. Consider the case where the firm finances the project with all equity. Doing so reduces the firm's net leverage ratio from 30% to 26.09%. To return its capital structure to its pre-project level of 30%, the firm can issue debt in the amount d or reduce its cash in the amount c by paying a dividend or repurchasing shares. How much?

$$\begin{aligned}0.30 &= \frac{(\text{Debt} + d) - \text{Cash}}{\text{Operating}} = \frac{(50 + d) - 20}{115} \implies d = 4.5 \\ 0.30 &= \frac{\text{Debt} - (\text{Cash} + c)}{\text{Operating}} = \frac{50 - (20 + c)}{115} \implies c = -4.5\end{aligned}$$

¹⁴See the study, “Do Firms Rebalance Their Capital Structures?” by Mark Leary and Michael R. Roberts in the 2005 Journal of Finance.

In other words, after financing the project with all equity, the firm can issue \$4.5 million of debt or reduce its cash balance by \$4.5 million to return its capital structure back to its pre-project level. If it does so, then the relevant project capital structure is a 30% net leverage ratio because this action is incremental; it only occurs because of the project.

Summary

Several lessons came from this exercise, which depend on the perfect capital markets assumption.

- Existing shareholders reap the rewards of positive NPV investments. The value of their shares increase by the NPV of the investment.
- It doesn't matter how we finance the investment because they all have the same implications for shareholder value.
- Issuing equity does not harm shareholders as long as the money raised is invested in positive NPV projects. So, equity dilution, in most instances, is a fallacy or, as we'll see below, limited to specific situations.
- Earnings per share will be higher with debt financing, all else equal, but the higher earnings are offset by an increase in risk that negates any value benefits. Put differently, shareholders should be ultimately concerned with share prices, not earnings per share.
- The financing choice has no impact on the operating cost of capital, or the value of the firm. It only impacts how the value is divided and the risk shared among debt and equity investors.
- Just like we focus on incremental cash flows, we also focus on incremental financing when assessing a project. Project financing may or may not change the capital structure of the firm, and therefore the equity cost of capital, depending on any actions the firm takes to offset a change.

11.5.3 Considering Market Imperfections

Transaction Costs

Transaction costs associated with issuing securities effectively increase the cost of the project and therefore reduce its NPV. For example, if an equity issuance costs 7% of the amount raised, then the total cost of the project including financing costs is $10 \times (1 + 0.07) = \$10.70$ million. This increased cost reduces the NPV of the project to $15 - 10.7 = \$4.3$ million. The process of quantifying the effect on shareholder (new and existing) value, EPS, price per share, and the P/E ratio is identical to what was done in the perfect capital markets case. In a nutshell, because the NPV of the project is reduced by financing costs, existing

shareholders benefit less than they would without the transaction cost - \$0.70 million in this example. Consequently, the stock price will go up by less than it would without transaction costs. New shareholders still buy shares at a fair price - zero NPV investment for them - and are no better or worse off. The story is the same for a debt issuance, though the numbers will differ depending on the financing costs.

Taxes

Accounting for corporate taxes has two immediate consequences for our analysis. First, our free cash flow is no longer equal to EBIT, it equals $EBIT \times (1 - \text{Tax rate})$ or net operating profit after taxes (NOPAT). (Recall we assumed that there are no working capital or long-term investment requirements.) Second, the interest from the debt creates a tax shield, which reduces payments to the government. Table 14 presents the implications of introducing corporate taxes into our project financing example when the effective and marginal tax rates are 21%.

Let's start by reassessing the firm before any project financing when taxes are present. To estimate enterprise value, we need to discount all the future NOPAT by the weighted average cost of capital (WACC) - equation ???. In other words,

$$\text{Enterprise value} = \frac{NOPAT}{r^{WACC}}$$

where

$$r^{WACC} = \frac{Debt - Cash}{\text{Enterprise value}} r^D (1 - \text{Tax Rate}) + \frac{Equity}{\text{Enterprise value}} r^E.$$

Recall that enterprise value is the same as operating asset value (*Operating*), which equals ($Equity + Debt - Cash$).

There are two challenges with estimating enterprise value. First, if our firm's capital structure changes, then so will the WACC. Remember figure 11.4. As leverage increases, the WACC decreases. So, if we want to discount all the future cash flows of the firm by the same number, we have to *assume* that the firm will maintain the same net leverage ratio, $(Debt - Cash) / \text{Enterprise value}$. Doing so, ensures that the WACC is constant over time absent a change to the business or the risk premia that investors demand. As suggested above, a constant leverage ratio is often not a bad approximation.

The second challenge is that in order to estimate the WACC, we need to know the enterprise value. But, to estimate the enterprise value we need to know the WACC. The two measures depend on one another. There are a couple of ways to tackle this. One way is to solve the two equations in terms of the variables we know.¹⁵ This approach only works

¹⁵The two equations to be solved are

$$\begin{aligned} \text{Enterprise value} &= \frac{NOPAT}{r^{WACC}} \text{ and,} \\ r^{WACC} &= r^O - \frac{Debt - Cash}{\text{Enterprise value} \times \text{TaxRate} \times r^D} \end{aligned}$$

	Financial Strategies					
	Pre-Issue	All Equity	All Debt	All Cash	50% Cash 50% Debt	45% Debt 55% Equity
P&L						
EBIT	10.00	11.50	11.50	11.50	11.50	11.50
Net interest expense	1.50	1.50	2.00	2.00	2.00	1.73
EBT	8.50	10.00	9.50	9.50	9.50	9.78
Taxes	1.79	2.10	2.00	2.00	2.00	2.05
Earnings	6.72	7.90	7.51	7.51	7.51	7.72
Assets - Cash						
Cash	20.00	20.00	20.00	10.00	15.00	20.00
Expected return on cash, r^C	5.00	5.00	5.00	5.00	5.00	5.00
Interest income, $r^C \times C$	1.00	1.00	1.00	0.50	0.75	1.00
Assets - Operating						
NOPAT, $EBIT \times (1 - \text{Tax Rate})$	7.90	9.09	9.09	9.09	9.09	9.09
Operating cost of capital, r^O	10.00	10.00	10.00	10.00	10.00	10.00
WACC, r^{WACC}	9.62	9.66	9.56	9.56	9.56	9.62
Enterprise value, $O = NOPAT/r^{WACC}$	82.15	94.00	95.05	95.05	95.05	94.47
Claims - Debt						
Debt value, D	50.00	50.00	60.00	50.00	55.00	54.50
Debt cost of capital, r^D (%)	5.00	5.00	5.00	5.00	5.00	5.00
Interest expense, $r^D \times D$	2.50	2.50	3.00	2.50	2.75	2.73
Claims - Equity						
Equity value, $E = V - D$	52.15	64.00	55.05	55.05	55.05	59.97
Equity cost of capital, r^E (%)	12.88	12.34	13.63	13.63	13.63	12.88
Earnings, Earn = EBIT - Interest	6.72	7.90	7.51	7.51	7.51	7.72
Shares outstanding (mil)	5.00	5.93	5.00	5.00	5.00	5.50
Price/share, P	10.43	10.80	11.01	11.01	11.01	10.89
Earnings per share, EPS=Earn/Shares	1.34	1.33	1.50	1.50	1.50	1.40
Price-to-earnings ratio, P/E=P/Earn	7.77	8.10	7.34	7.34	7.34	7.77
Firm						
Firm value, $V = C + O$	102.15	114.00	115.05	105.05	110.05	114.47
Asset cost of capital, r^A	9.02	9.12	9.13	9.52	9.32	9.13
Unlevered firm value, $NOPAT/r^O + C$	99.00	110.85	110.85	100.85	105.85	110.85
Value interest tax shield	3.15	3.15	4.20	4.20	4.20	3.62

Table 14: Pre- and Post-Project-Financing Information with Corporate Taxes. All Dollar Values are in Millions Except Per Share Figures

Their solution is

$$\text{Enterprise value} = \frac{NOPAT}{r^O} \times \left(1 + \frac{Debt - Cash}{NOPAT} \times TaxRate \times r^D \right)$$

$$r^{WACC} = \frac{1}{1 + \frac{Debt - Cash}{NOPAT} \times TaxRate \times r^D}$$

Copyright © 2020 Michael R. Roberts

in our simplified example in which the firm (and project) are perpetuities. In practice, this approach isn't practical. The alternative is to let Excel figure it out, a process detailed in the Circular Reference section of the Spreadsheet Appendix

The value of the firm is the sum of cash and enterprise value, $20+82.15=\$102.15$ million, or equivalently debt and equity. Because debt is assumed to be \$50 million, equity is $102.15-50=\$52.15$ million. The equity cost of capital is

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.10 + \frac{50 - 20}{52.15} (0.10 - 0.05) = 12.88\%$$

Earnings from the P&L are \$6.72 million, which when discounted by the equity cost of capital recovers the equity value, $6.72/0.1288=\$52.15$ million. Price and earnings per share follow by dividing the equity value by the number of shares outstanding.

The asset cost of capital, r^A , which captures both operating assets and cash is

$$\begin{aligned} r^A &= \frac{Cash}{Value} r^C + \frac{Operating}{Value} r^O = \frac{20}{102.15} \times 0.05 + \frac{82.15}{102.15} \times 0.10 = 9.02\%, \text{ or} \\ r^A &= \frac{Debt}{Value} r^D + \frac{Equity}{Value} r^E = \frac{50}{102.15} \times 0.05 + \frac{52.15}{102.15} \times 0.1288 = 9.02\% \end{aligned}$$

This estimate is lower than the operating cost of capital, $r^O = 10\%$, because of the low risk cash earning 5%.

We saw that the enterprise value of the firms is \$82.15 million, which we computed by discounting NOPAT by the WACC. Because of the interest tax shield, this value embeds two components: (1) the enterprise value as if it were unlevered and (2) the value of the interest tax shield. We can estimate these two components separately to determine how much value financing is bringing to our business. We just have to discount the NOPAT by the operating cost of capital, or **unlevered cost of capital**, r^O .

$$\text{Unlevered enterprise value} = \frac{NOPAT}{r^O} = \frac{7.90}{0.10} = 79.00$$

The difference between the enterprise value and unlevered enterprise value is the value of the interest tax shield, $102.15 - 79.00 = \$23.15$ million. This value can also be computed by discounting the periodic interest tax shield by the operating cost of capital. Every period the firms saves in taxes its net interest expense times its tax rate, or $(2.50-1.00) \times 0.21 = 0.315$. The present value of this perpetuity at 10% is \$3.15 million. Note, we discount the interest tax shield by the operating cost of capital because we have assumed the firm will maintain a constant net leverage ratio.

People often talk about the **levered** and **unlevered value of the firm**, as opposed to enterprise value. This just requires adding the excess cash to enterprise value.

$$\begin{aligned} \text{Levered firm value} &= \frac{NOPAT}{r^{WACC}} + Cash = \frac{7.90}{0.0926} + 20 = \$102.15 \\ \text{Unlevered firm value} &= \frac{NOPAT}{r^O} + Cash = \frac{7.90}{0.10} + 20 = \$99.00 \end{aligned}$$

The difference between the levered and unlevered value of the firm, like enterprise value, is equal to the value of the interest tax shield, \$3.15 million.

Now let's consider the implications of the different financing strategies beginning with all equity financing. The project NOPAT is now $1.50 \times (1 - 0.21) = \1.19 million per year. The project WACC assuming all equity financing is equal to the operating cost of capital, r^O . Recall the WACC from equation 11.13

$$\text{Project WACC} = r^O - \frac{\text{Debt} - \text{Cash}}{\text{Equity} + \text{Debt} - \text{Cash}} \times \text{TaxRate} \times r^D$$

No debt ($D = 0$) and no excess cash ($C = 0$) mean the net leverage ratio equals zero. Thus, the WACC with all equity financing is just the operating cost of capital. The key assumption here is that the incremental financing is truly all equity. In other words, the firm won't make any other changes to its capital structure (e.g., reducing cash or issuing debt) in response to the project financing.

The value of the project is $1.19/0.10 = \$11.85$ million, and the NPV is $11.85 - 10.00 = \$1.85$ million. If our intuition from the baseline case is correct, the enterprise and equity value should both increase by the value of the project, the latter because equity is footing the entire bill. Let's verify these statements from first principles using our fundamental value relation. The NOPAT of the firm with the project is $11.50 \times (1 - 0.21) = \9.09 million. We need to discount this by the *company*, as opposed to *project*, WACC. But to estimate the project, we need to know the enterprise value of the firm. So, we're in a similar situation as before where we need to solve for these two quantities at the same time. Fortunately, spreadsheets easily handle this situation and we get a WACC of 9.66% and an enterprise value of \$94.00 million. Let's double check these values.

$$r^{WACC} = 0.10 + \frac{50 - 20}{94} \times 0.21 \times 0.05 = 9.66\%$$

$$\text{Enterprise value} = \frac{9.09}{0.0966} = \$94.00 \text{ million}$$

Equity equals the enterprise value minus debt plus cash or $94.00 - 50.00 - 20.00 = \64.00 million. We can verify this number by discounting the equity cash flows by the equity cost of capital. But, to find the equity cost of capital, we need to know the value of equity per equation 11.5. Again, the spreadsheet comes to the rescue.

$$r^E = 0.10 + \frac{50 - 20}{64} (0.10 - 0.05) = 12.34\%$$

$$\text{Equity value} = \frac{7.90}{0.1234} = 64.00$$

The question now is: How many shares have to be issued and at what price to finance the project? We know the following two equalities must hold.

$$\text{Shares}_{new} \times \text{Price} = 10$$

$$(\text{Shares}_{old} + \text{Shares}_{new}) \times \text{Price} = 64$$

The first equation says that the number of new shares we issue times the price equals the cost of the project, \$10 million. The second equation says that the total number of shares times the price must equal the new equity value, \$64 million. Solving these two equations shows that the firm must issue 0.93 million shares at \$10.80 per share.

The value of the firm is just the sum of cash and enterprise value, or equivalently debt and equity. So, our firm is worth $94 + 20 = 50 + 64.00 = \$114.00$ million. Similarly, the asset cost of capital is a weighted average of the cash and operating asset, or debt and equity, cost of capitals.

$$r^A = \frac{20}{20 + 94}0.05 + \frac{94}{20 + 94}0.10 = 9.12\%$$

$$r^A = \frac{50}{50 + 64}0.05 + \frac{64}{50 + 64}0.1234 = 9.12\%$$

The asset cost of capital has increased from 9.02% to 9.12% because of the increase in risky operating assets relative to safe cash.

If we finance the project with all debt, we need to let the spreadsheet figure out the project WACC and value, which depend on one another through the equity value, E .

$$\text{Project WACC} = -\frac{\text{Debt} - \text{Cash}}{\text{Equity}} \times \text{TaxRate} \times r^D = 0.10 - 1 \times 0.21 \times 0.05 = 9.19\%$$

$$\text{Project value} = \frac{\text{NOPAT}}{r^{\text{WACC}}} = \frac{1.19}{0.0919} = \$12.90 \text{ million}$$

The project NPV is $12.90 - 10.00 = \$2.90$, larger than the NPV of the project with all equity financing (\$1.85 million) because of the interest tax shield which is worth $2.90 - 1.85 = \$1.05$ million.

Looking at the firm, debt increases by \$10 million to fund the project, while equity increases by the NPV of the project, \$2.90 million. The value of the firm is the sum of debt and equity, $60.00 + 55.05 = \$115.05$ million. Enterprise value is equity plus debt minus cash, $115.05 - 20.00 = \$95.05$ million. We can (and should) verify all of these valuations using the fundamental value relation.

The enterprise WACC from equation 11.13 is

$$r^{\text{WACC}} = 0.10 - \frac{60 - 20}{95.05} \times 0.21 \times 0.05 = 9.56\%.$$

Discounting the firm's NOPAT by the enterprise WACC recovers the enterprise value, $9.09/0.0956 = \$95.05$ million. The equity cost of capital from equation 11.5 is

$$r^E = 0.10 + \frac{60 - 20}{55.05} (0.10 - 0.05) = 13.63\%.$$

Discounting the firm's earnings by the equity cost of capital recovers the equity value, $7.51/0.1363 = \$55.05$ million. The new share price, $55.05/5 = \$11.01$, is higher than the all equity financing share price (\$10.80) because of the interest tax shield.

Looking at the remaining financing options reveals results analogous to what we found in the perfect markets setting. All cash and mixed cash-debt financing produces results identical to the all debt financing scenario from the perspective of shareholders, though the overall firm is smaller and riskier because of the reduction in cash. The 45-55 debt equity mix produces results in between the all equity and all debt financing outcome and maintains a constant WACC and equity cost of capital. Overall, more debt financing leads to greater value for shareholders because of the interest tax shield as seen in higher stock prices and the higher interest tax shield.

11.6 Project Financing's Impact on the Project

The previous section illustrated how different project financing impacts the firm and, to a lesser extent, the project. This section emphasizes the project and considers the effect of different financial strategies on the cost of capital, project valuation, and the approach used to value the project. To emphasize the financing, we'll use the Dell tablet case study from chapter 5. Only now, rather than assuming that the relevant cost of capital for the project is 12%, we'll compute the relevant cost of capital from first principles and show how it varies with different financial strategies.

11.6.1 Dell's MVBS and The Tablet Project

Recall that the date is January 2012 and Dell is considering introducing a tablet to compete with Apple's iPad. (This is hypothetical.) Dell's market value balance sheet containing both values and expected returns/cost of capitals at this time is presented in figure 15.

Assets		Claims	
Cash	14,818	Debt	9,254
Expected return on cash, r^C	3.09	Debt cost of capital, r^D	3.09
Operating	25,390	Equity	30,954
Expected return on operating assets, r^O	7.81	Equity cost of capital, r^E	6.96
Total assets	40,211	Total claims	40,211
Expected return on assets, r^A	6.07	Asset cost of capital, r^A	6.07

Table 15: Dell Inc.'s Market Value Balance Sheet as of January 2012 (\$mil)

The project P&L and free cash flow schedule are shown in Figures 11.13 and 11.14. (Refer to chapter 5 for details on the construction of these figures.) The marginal tax rate is assumed to be 21%.

P&L (\$mil)	0	1	2	3
Sales, tablet	\$0.0	\$80.0	\$1,100.0	\$4,356.0
Cannibalization, laptop sales	0.0	(5.0)	(10.0)	(20.0)
Net sales	0.0	75.0	1,090.0	4,336.0
COGS	0.0	(15.0)	(193.1)	(716.1)
Gross profit	0.0	60.0	896.9	3,619.9
SG&A	0.0	(48.0)	(660.0)	(2,613.6)
R&D	(50.0)	(10.0)	(10.0)	(10.0)
EBITDA	(50.0)	2.0	226.9	996.3
Depreciation and amortization		(39.0)	(39.0)	(39.0)
EBIT	(50.0)	(37.0)	187.9	957.3
Taxes	10.5	7.8	(39.5)	(201.0)
NOPAT	(39.5)	(29.2)	148.4	756.3

Figure 11.13: Dell Tablet - P&L

Free cash flow (\$mil)	0	1	2	3
NOPAT	(39.5)	(29.2)	148.4	756.3
Depreciation and amortization	0.0	39.0	39.0	39.0
NLTI	(400.0)			217.4
NWCI	(3.8)	(45.6)	(148.0)	197.3
Unlevered free cash flow	(443.3)	(35.8)	39.5	1,210.0

Figure 11.14: Dell Tablet - Free Cash Flows

11.6.2 Target Capital Structure & WACC Valuation

Let's assume that Dell would like to maintain a 40% project leverage ratio. What is the WACC? How much debt should be allocated to the project each period? What is the value of the project? How much of this value comes from the financial policy versus the project? These are the questions we'd like to answer.

With a target or constant leverage ratio, the easiest way to value the project is by computing the after-tax weighted average cost of capital (equation 11.13) and discounting the unlevered free cash flows. This valuation approach is referred to as **WACC valuation**. Assuming a constant leverage ratio, simplifies two challenges. First, we don't need to know the value of the project to compute the WACC because we are assuming a value for the (net) leverage ratio, $Debt/(Equity + Debt)$ or $(Debt - Cash)/(Equity + Debt - Cash)$. In

our example, both measures are assumed to be equal to 40% every period. Second, because these ratios are assumed to be constant, the WACC doesn't change over time. We can discount each unlevered cash flow by the same number.

The project WACC assuming a 40% leverage ratio equals

$$\begin{aligned} r^{WACC} &= r^O - \frac{Debt - Cash}{Equity + Debt - Cash} r^D (1 - TaxRate) \\ &= 7.81 - 0.40 \times 3.09 \times 0.21 \\ &= 7.55\%. \end{aligned}$$

The NPV of the project is therefore

$$NPV = -443.3 + \frac{35.8}{1 + 0.0755} + \frac{39.5}{(1 + 0.0755)^2} + \frac{1,210.0}{(1 + 0.0755)^3} = 530.10.$$

If we discount the project by the operating (a.k.a., unlevered) cost of capital, r^O , we get the *unlevered* NPV of the project.

$$Unlevered\ NPV = -443.3 + \frac{35.8}{1 + 0.0781} + \frac{39.5}{(1 + 0.0781)^2} + \frac{1,210.0}{(1 + 0.0781)^3} = 523.02$$

This is the value of the project had we financed it entirely with equity and received no interest tax shields. The difference between the NPV using the WACC and that using the operating cost of capital is present value of all the interest tax shields over the life of the project, i.e., the value of the financial policy - $530.10 - 523.02 = \$7.08$ million.

But, what exactly is the debt policy of the firm? In other words, how much debt does the firm have to issue or retire every period over the life of the project? To answer this, we need to know the value of the project every period because our assumption is that 40% of the value is funded by debt. But, our fundamental value relation tells us that the financial value of any asset or opportunity is just the sum of the discounted future cash flows.

$$\begin{aligned} Value_3 &= \frac{0.0}{1 + 0.0755} = \$0.0 \implies Debt_3 = 0.40 \times 0.0 = \$0.0 \text{ million} \\ Value_2 &= \frac{1,210.0}{1 + 0.0755} = \$1,125.0 \implies Debt_2 = 0.40 \times 1,125.0 = \$450.0 \text{ million} \\ Value_1 &= \frac{(39.5 + 1,125.0)}{1 + 0.0755} = \$1,082.7 \implies Debt_1 = 0.40 \times 1,082.7 = \$433.1 \text{ million} \\ Value_0 &= \frac{(-35.8 + 1,082.7)}{1 + 0.0755} = \$973.4 \implies Debt_0 = 0.40 \times 973.4 = \$389.34 \text{ million} \end{aligned}$$

The value relations above may look new but each is just our fundamental value relation written in a more compact way. There are no cash flows after year three so the value of the project in year three is zero. In year two, there is one more cash flow in year three, which we discount by the WACC to get a value of \$1,125.0. In year 1, there are two cash flows, which we can discount as we usually do.

$$Value_1 = \frac{39.5}{1 + 0.0755} + \frac{1,210.0}{(1 + 0.0755)^2} = \$1,082.7 \text{ million}$$

But, because $Value_2 = 1,210.0/(1 + 0.0755)$, we can rewrite the value one year from today as the discounted sum of year two's cash flow and value.

$$Value_1 = \frac{39.5}{1 + 0.0755} + \underbrace{\frac{1,210.0}{1 + 0.0755}}_{Value_2} \times \frac{1}{1 + 0.0755} = \frac{39.5 + Value_2}{1 + 0.0755} = \$1,082.7 \text{ million}$$

Similarly, we can rewrite the project value today as the discounted sum of year one's cash flow and value.

$$\begin{aligned} Value_0 &= \frac{35.8}{1 + 0.0755} + \frac{39.5}{(1 + 0.0755)^2} + \frac{1,210.0}{(1 + 0.0755)^3} \\ &= \frac{35.8}{1 + 0.0755} + \underbrace{\left(\frac{39.5}{(1 + 0.0755)} + \frac{1,210.0}{(1 + 0.0755)^2} \right)}_{Value_1} \times \frac{1}{1 + 0.0755} \\ &= \frac{35.8 + Value_1}{1 + 0.0755} \end{aligned}$$

In general, the value at any time t is equal to the discounted sum of next period's cash flow and value.

$$Value_t = \frac{CashFlow_{t+1} + Value_{t+1}}{1 + r} \quad (11.14)$$

This result isn't new. We saw it in chapter 8 when we introduced the dividend discount model, which shows that the price of a stock in period t is equal to the dividend (cash flow) plus the price (value) next period.

Figure 11.15 summarizes the valuation results. The two results we have yet to discuss are the equity value and equity cost of capital. The former follows from our value and debt estimates because equity must equal value minus debt according to our market value balance sheet which holds for firms and projects. The equity cost of capital follows from equation 11.5 and recognizing that there is no excess cash associated with the project.

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 7.81 + \frac{0.40}{1 - 0.40} (7.81 - 3.09) = 10.96\%$$

Notice that the equity cost of capital is constant over the life of the project because the capital structure is constant - a fact we can verify by dividing the debt of the project by the value of the project each period.

11.6.3 Predetermined Capital Structure & Adjusted Present Value

Now let's examine the project when the financial policy is known in advance or **predetermined**. For example, let's assume that the debt policy for the project is to borrow \$400 million today at 3.09% and then pay down the principal in equal amounts each year over the life of the project as show in Table 16

WACC Valuation (\$mil)	0	1	2	3
WACC	7.55%	7.55%	7.55%	7.55%
Project value	973.35	1,082.67	1,124.99	0.00
Project NPV	530.10			
Check	530.10			
Debt schedule				
Debt	389.34	433.07	450.00	0.00
Change in debt	389.34	43.73	16.93	(450.00)
Interest	0.00	12.01	13.36	13.88
ITS	0.00	2.52	2.81	2.92
ITS value	7.08	5.11	2.70	0.00
Equity value	584.01	649.60	674.99	0.00
Equity cost of capital	10.96%	10.96%	10.96%	10.96%
Equity value check	584.01	649.60	674.99	0.00
Project NPV check	530.10			
Unlevered value	966.27	1,077.56	1,122.28	0.00
Interest tax shield value	7.08	5.11	2.70	0.00

Figure 11.15: Dell Tablet - WACC Valuation with 40% Leverage Target

	0	1	2	3
Debt	400	266.7	133.3	0.0

Table 16: Dell Tablet - Predetermined Debt Policy (\$mil)

The difficulty in using the WACC valuation approach in this setting is that we don't know what leverage ratio corresponds to the this debt policy because we don't know the value of the project. But, to estimate the value of the project, we need to know the WACC. We saw this circularity issue above, which is easily handled in Excel. However, there is another valuation approach that avoids this complication when capital structure is predetermined - the **adjusted present value** or **APV** method. The APV is a three step process.

1. Estimate the unlevered value of the project by discounting the unlevered free cash flows by the unlevered (i.e., operating) cost of capital.
2. Estimate the value of the interest tax shield by discounting the interest tax shields by the debt cost of capital.
3. Add 1. and 2.

Figure 11.16 presents the results of this process.

The unlevered value of the project in period zero is

$$Unlevered\ Value_0 = \frac{-35.8}{1 + 0.0781} + \frac{39.5}{(1 + 0.0781)^2} + \frac{1,210.0}{(1 + 0.0781)^3} = \$966.3 \text{ million.}$$

APV Valuation (\$mil)	0	1	2	3
Debt schedule				
Debt	400.00	266.67	133.33	0.00
Change in debt	400.00	(133.33)	(133.33)	(133.33)
Interest	0.00	12.34	8.23	4.11
ITS	0.00	2.59	1.73	0.86
ITS value	4.93	2.49	0.84	0.00
Unlevered value	966.27	1,077.56	1,122.28	
Project value	971.20	1,080.05	1,123.12	
Equity value	571.20	813.39	989.79	
Project NPV	527.95			

Figure 11.16: Dell Tablet - WACC Valuation with 40% Leverage Target

To compute the value of the interest tax shield, we first need to estimate the interest and corresponding tax shield period by period. For example, the interest paid in year one is the cost of debt times the debt balance as of the previous period, $400 \times 0.0309 = \$12.34$ million. The interest tax shield (ITS) is the product of interest expense and the tax rate, $12.34 \times 0.21 = \$2.59$ million. Because of this project's financing choice, Dell will pay \$2.59 million less in taxes one year from today assuming they do nothing to offset this financing (e.g., issue equity, retire other debt). The value of all the future interest tax shields created by the project financing is

$$\text{Interest Tax Shield Value}_0 = \frac{2.59}{1 + 0.0309} + \frac{1.73}{(1 + 0.0309)^2} + \frac{0.86}{(1 + 0.0309)^3} = \$4.93 \text{ million}$$

Adding the unlevered value of the project and the value of the interest tax shield gives us the levered value of the project, $966.3 + 4.93 = \$971.20$ million. Coupled with the debt schedule, we can back out the equity value as the difference between the project and debt values.

11.6.4 Summary

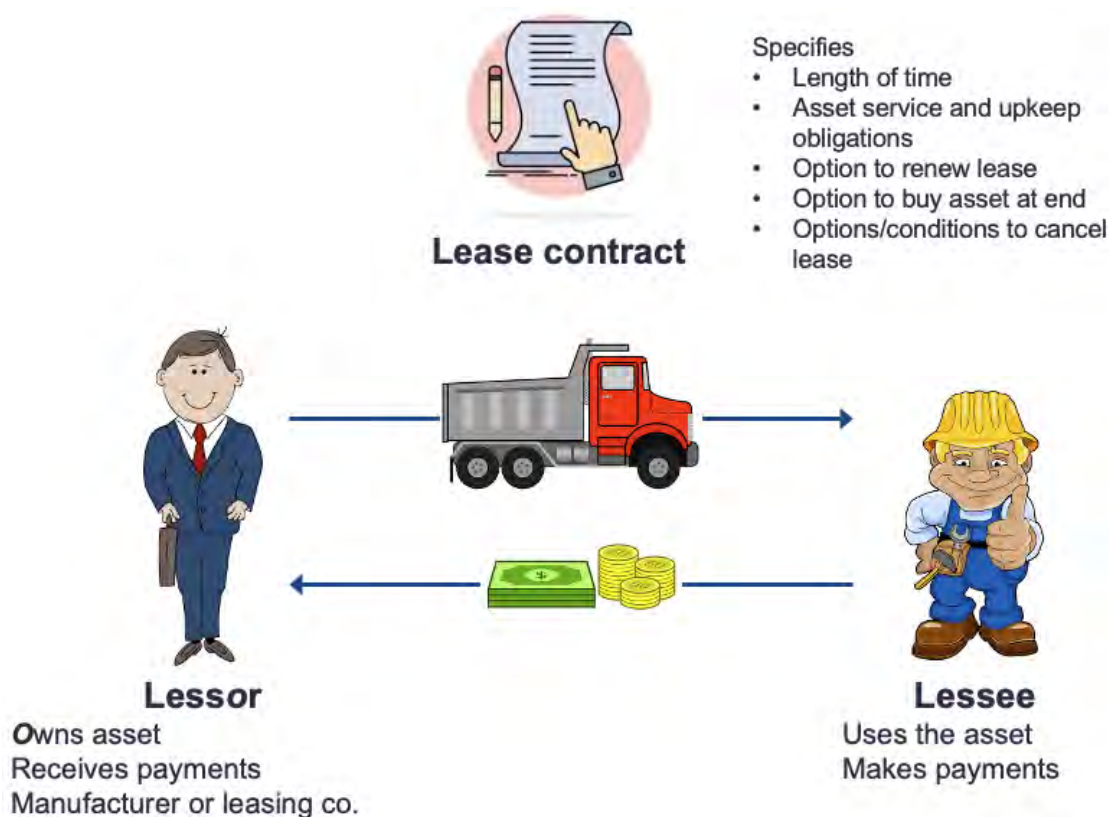
When project leverage targets a fixed level (e.g., 40%) over the life of the project, WACC valuation is straightforward. Compute the WACC using the target leverage ratio, discount the unlevered cash flows by the WACC, and sum. When the debt policy is predetermined and project leverage not necessarily constant, APV valuation is straightforward. Estimate the unlevered value of the project using the unlevered cost of capital, estimate the value of the interest tax shield using the debt cost of capital, and sum.

We might ask: Can we use APV when we have target leverage ratio, or WACC when we have a predetermined debt policy? Yes, and we'll get the exact same answer. In other words, we can use any valuation method we want. When done correctly, they all generate the same answer. Some are simply more convenient to use in certain situations.

Consider the target capital structure scenario above. Figure 11.15 shows the debt policy implied by a constant leverage ratio of 40%. It also shows the value of the interest tax shield, \$7.08 million. If we add this to the unlevered value of the project, \$966.27 million, we get \$973.35 million, which is the same value we got when we did the WACC valuation. In the predetermined debt scenario, we can perform a WACC valuation, but doing so requires a modification to our after-tax WACC formulae, which we relegate to the technical appendix.

11.7 Leasing

A **lease** is a contract between two parties - a **lessor** and a **lessee**. The lessor owns an asset that gives to the lessee for use over a finite amount of time in exchange for periodic payments. (I keep the two names - lessor and lessee - straight by remembering that the “o” in lessor refers to “owner.”) This arrangement is illustrated in Figure 11.17. The lessor owns a truck that it lends to the lessee for some period of time. During that time, the lessee makes periodic (often monthly) payments to the lessor for the right to use the lessor’s truck. The lease contract specifies a number of things including the term of the arrangement, the lease payments, the obligations of both parties regarding maintenance and upkeep of the asset, as well as insurance, and options to purchase the asset, cancel the contract, and renew the contract.



For example, when we lease a car from a car dealership we agree to make monthly payments to the dealership over some time period - typically two to seven years - in exchange for using the car during that period. While using the car, we agree to insure and maintain the car in good working order. At the end of the lease term, we have to return the car or we can purchase the car for a previously agreed upon price. The essential difference from purchasing the car by borrowing money is that in a lease we don't own the car, it belongs to the lessor. This difference has some interesting and important implications for decision making that we're going to explore in this section. Before doing so, let's understand the different varieties of leases and their treatment by accountants and the tax authority, all of which can affect the decision making surrounding leases.

11.7.1 Different Flavors of Leases

There are many different types of leases some distinguished by the lessor. A **direct lease** is a lease in which the lessor is an independent leasing company whose primary business is acquiring assets for the purpose of leasing them out. Some examples include United Rentals, which leases industrial and construction equipment, BOC Aviation, which leases aircraft, and Ryder, which leases trucks, vans, and trailers. A **sale-type lease** is a lease in which the lessor is the manufacturer of the asset, such as an automaker or aircraft company. The Ford Motor Company and Boeing are two examples of car and aircraft manufacturers that lease many of their products. A **sale-leaseback** is a lease in which the owner of the asset sells the asset and then leases it back from its new owner. For example, the household retailer Bed, Bath, and Beyond sold 2.1 million square feet of commercial real estate to Oak Street in 2019 and then leased this real estate back from Oak Street. Another example is when I - yes, the author - sold our home and then leased it back from the new owner while my family searched for a new home. The motivation in both instances was, in part, an injection of cash for other purpose such as buying a new home. Finally, a **leveraged lease** is a lease in which the lessor borrows a significant sum of money to acquire the asset, often 50% or more, and then leases the asset as a means to repay the debt and earn a profit. Leveraged leases can get substantially more complicated from a legal perspective with the use of **special purpose entities** or **SPEs**.

Assets typically require maintenance and upkeep. Many assets come with recurring fees such as registration or certification. Sufficiently expensive assets typically need to be insured. Leases often differ in terms of who bears responsibility for these tasks. In **full-service** or **rental** leases, the lessor is responsible for all asset-related tasks. Some commercial real estate leases are full-service in that the landlord will pay for all operating expenses for the property. A **net lease** has the lessee responsible for the tasks. For example, most auto leases are net leases in which the lessee is responsible for maintaining, insuring, and registering the car. However, many leases are a combination of full-service and net allocating responsibilities to both lessor and lessee.

Leases also differ in how they end. In some instances, the lessee simply returns the asset to the lessor and that's the end of it. Other leases will offer an option to renew or extend the lease for another period. However, many leases will offer the lessee the option to purchase

the asset sometime during - **early buyout** - or at the end of the lease. A **lease-end buyout** (a.k.a., **residual, fixed-price**) is perhaps the most common option in which the lessee can purchase the asset for its residual value established at the start of the lease. A **market value buyout** let's the lessee purchase the asset at the end of the lease at the prevailing market price. A **bargain purchase option** let's the lessee purchase the asset at a price significantly below market value or at a negligible price (e.g., \$1.00) at the end of the lease.

With a purchase option, the lease acts as a trial period during which the lessee can decide whether they want to keep the asset longer-term. While we won't value these options, which require option valuation tools beyond the scope of this text, they can play an important role in the determination of lease payments and the leasing decision. Therefore, it's useful to think of our discussion and analysis here as providing the foundation for analyzing leasing decisions.

11.7.2 A Commercial Truck Lease

Let's set up an example to explore the leasing decision. Imagine we run a construction company in need of a large dump truck, which costs \$120,000, has a 10-year depreciable life after which it is worthless. Also assume that the operating and debt cost of capitals for our company are 12% and 6%, respectively, and that our tax rate is 21%. Should we buy the truck or lease it for five years if the residual value at the end of the lease is \$60,000? Obviously, this choice will depend on the lease payments - size and timing. It will also depend on how the lease is classified.

The **Financial Accounting Standards Board**, or **FASB**, distinguishes between two type of lease - **operating** and **financial** (or **capital**) - using the following set of criteria. A lease is deemed a financial lease if one or more the following conditions are met.¹⁶

1. The lease transfers ownership of the asset to the lessee by the end of the lease term.
2. There is a high likelihood that the lessee will exercise a purchase option at the end of the lease.
3. The lease term is for at least 75% of the assets economic life, assuming the least does not begin at or near the end of the asset's economic life.
4. The sum of the present values of the lease payments is less than 90% of the market value of the asset.
5. The asset is so specialized that it has no alternative use to the lessor at the end of the lease.

If none of these conditions are met, then the lease is deemed an operating lease. In our example, none of these conditions are met so the lease would be classified as an operating

¹⁶These conditions are based on Accounting Standards Codification (ASC) 842, "Leases."

lease. More than just a label, we'll see that operating and financial leases have different accounting treatments that impact the taxes the lessor and lessee pay. However, before exploring this facet of leasing, let's start with the perfect capital markets benchmark.

11.7.3 The Perfect Capital Markets Benchmark

What must the truck lease payments be when capital markets are perfect? Whatever makes shareholders indifferent between buying the truck and leasing it. If not, then leasing could create or destroy value for shareholders. If we buy the truck, we'll spend \$120,000 today. If we lease the truck, we will have to make five annual payments at the *start* of each year. After five years, we'll return the truck, which has an expected market or residual value equal to \$60,000. These cash flows, all of which are outflows, are illustrated in Figure 11.18. Let "LP" denote lease payment.

Years	0	1	2	3	4	5
	----- ----- ----- ----- -----					
Lease	LP	LP	LP	LP	LP	Residual
Buy	120					

Figure 11.18: Truck Lease Cash Flows (\$000s) in Perfect Capital Markets

Equivalent between the two options means the cost of buying the truck equals the cost of leasing the truck.

$$\underbrace{120,000}_{\text{Cost to buy}} = \underbrace{LP + \frac{LP}{1+0.06} + \frac{LP}{(1+0.06)^2} + \frac{LP}{(1+0.06)^3} + \frac{LP}{(1+0.06)^4} + \frac{60,000}{(1+0.06)^5}}_{\text{Cost to lease}}$$

Solving for the lease payment, LP , yields \$16,833.76.¹⁷ Discounting the future cash flows by the debt cost of capital, 6%, is reasonable for the lease payments (LP), which are fixed obligations similar to payments on senior secured credit. Discounting the residual value by the cost of debt is implausible. There is much greater risk associated with this cash flow so the operating cost of capital makes more sense. In this case, the lease payments are \$19,250.25, which are found by solving the following equation for LP .

$$120,000 = LP + \frac{LP}{1+0.06} + \frac{LP}{(1+0.06)^2} + \frac{LP}{(1+0.06)^3} + \frac{LP}{(1+0.06)^4} + \frac{60,000}{(1+0.12)^5}$$

¹⁷(The lease payments are an annuity starting today. We can use the present value of an annuity result (equation 2.1 modified to recognize that the payments start today.

$$PV = CF + \frac{CF}{r} \left(1 - (1+r)^{-(T-1)}\right)$$

The cash flow corresponding to this annuity is

$$CF = \frac{PV}{1 + \frac{1}{r} \left(1 - (1+r)^{-(T-1)}\right)}$$

The lease payments are higher using the higher discount rate because the present value of the truck is less. In other words, we're effectively using up more of the truck's value when the discount rate is higher or, equivalently, when the residual value is lower.

The message of this exercise is that when capital markets are perfect, the lease payment is uniquely determined. For example, imagine the lease payments were less than \$19,250.25. Demand for leases would increase, driving up the price until it reached \$19,250.25. Similarly for lease payments greater than \$19,250.25. More broadly, when capital markets are perfect, the lease decision is irrelevant because the firm is indifferent between leasing and buying by borrowing.

11.7.4 Considering Market Imperfections

Let's focus our attention on taxes, which are the primary, but only, distinction from borrowing.¹⁸ Should we buy or lease the truck when we face corporate taxes?

Buy the Truck

If we buy the truck, we have to pay the purchase price, \$120,000, but we'll get a depreciation tax shield. Assuming the truck is depreciated on a straight-line basis over its 10-year life, this means \$12,000 of depreciation per year and a tax shield equal to $12,000 \times 0.21 = \$2,520$ per year. Finally, at the end of five years, we'll have a truck that we can sell for \$60,000. (Remember our discussion of project viability in chapter 5. Assets and liabilities do not just disappear.) The cash flows are detailed in Figure 11.19.

Years	0	0	0	0	0	0
CapEx	-120					
Dep. tax shield		2.52	2.52	2.52	2.52	2.52
Residual value						60.0

Figure 11.19: Buy the Truck Cash Flows (\$000s)

We're not going to sum these cash flows period-by-period and then discount because they are not all the same risk. The depreciation tax shield has a different risk profile than the residual value, which means they should have different discount rates. Let's start with the depreciation tax shield. As long as the firm anticipates generating enough taxable income over the lease horizon - five years in this example - to utilize the depreciation tax shield, these cash flows may be considered relatively safe. Therefore, the debt cost of capital is a natural choice to discount these cash flows, but it's the wrong choice.

¹⁸Leases deemed **true leases** as opposed to **security interests** come with even higher priority than secured creditors in a firm's capital structure. In default, lessors of true leases can demand lessees continue their payments or seize their leased assets. This advantage may make leasing cheaper than certain forms of credit.

We want to discount these cash flows by the *after-tax* cost of debt, $r^D(1 - TaxRate) = 0.06 \times (1 - 0.21) = 4.74\%$. Why? Think about the purchase decision from the perspective of a project, which it is. The WACC defined earlier in equation 11.13 is

$$r^{WACC} = r^O - \frac{Debt - Cash}{Equity + Debt - Cash} \times TaxRate \times r^D.$$

(This is a project so there is no excess cash for the project, $C = 0$.) If we buy the truck, we'll finance it entirely with debt implying that the debt-to-value ratio for the project is 100%. As we just discussed, the risk of these cash flows assuming continued profitability is best represented by r^D , not r^A (or r^O). So, we can replace r^O in the equation above with r^D . Therefore, the appropriate WACC for the depreciation tax shield is just the after-tax cost of debt.

$$r^{WACC} = r^D - TaxRate \times r^D = r^D(1 - TaxRate) = 0.06 \times (1 - 0.21) = 4.74\%$$

The present value of the depreciation tax shields can be found using our annuity result (equation 2.1).

$$\frac{2,520}{0.0474} \times (1 - (1 + 0.0474)^{-5}) = \$10,989.14.$$

The residual value we said earlier has a risk profile more like the firm itself and therefore its risk is better represented by the operating cost of capital, r^O . This is just an assumption like using r^D as a basis for the depreciation tax shield. With a risk profile represented by the operating cost of capital, the appropriate discount rate for the residual value is the WACC from equation 11.13, again with no excess cash and 100% debt financing.

$$\begin{aligned} r^{WACC} &= r^O - \frac{Debt - Cash}{Equity + Debt - Cash} \times TaxRate \times r^D \\ &= 0.12 - 1.0 \times 0.21 \times 0.06 = 10.74\% \end{aligned}$$

The present value of the residual value is therefore

$$\frac{60,000}{(1 + 0.1074)^5} = \$36,027.05.$$

Now we can add the present values of (1) the cost of the truck, (2) the depreciation tax shield, and (3) the residual value to get the NPV of buying the truck by borrowing money.

$$-120,000 + 10,989.14 + 36,027.05 = -\$72,983.81$$

This is the total cost in today's dollars of purchasing the truck considering only the cash flows that are incremental to leasing. That is, this number doesn't mean that buying the truck is a negative NPV investment. We haven't considered any other implications such as revenue from purchasing the truck because we're assuming they're the same regardless of whether we buy or lease the truck. Here, we're focused exclusively on cash flows that differ from the leasing decision, which we explore next.

Lease the Truck

If we lease the truck, we make five annual lease payments at the start of each year. Because the lease is an operating lease, the payments are considered rental expenses and are fully tax deductible. (They are usually part of Selling, General, and Administrative expenses on the income statement.) The annual lease payment is \$19,250.25, so the after-tax cost of the lease each year is $19,250.25 \times (1 - 0.21) = \$15,207.70$. The cash flows associated with leasing are illustrated in Figure 11.20.




Figure 11.20: Lease the Truck Cash Flows (\$000s)

The lease payments are fixed obligations similar to debt payments so the risk is represented by the debt cost of capital, r^D . However, as with the buy option, the lease payments are equivalent to 100% debt financing and therefore the appropriate discount rate is *after-tax* cost of debt, $r^D(1 - TaxRate) = 0.06 \times (1 - 0.21) = 4.74\%$. Recognizing that the lease payments correspond to an annuity starting immediately, we can use equation 2.1 to value the last four payments and add the one occurring immediately.

$$15,207.70 + \frac{15,207.70}{0.0474} (1 - (1 + 0.0474)^{-4}) = \$69,460.72$$

Because the lease payments are outflow, this value should be negated, i.e., $-\$69,460.72$.

The Decision

Now that we know the value of the buy and lease decisions, we can simply choose the most valuable (least costly) option, which is leasing. Leasing saves us $-\$69,460.72 - (-\$72,983.81) = \$3,523.09$ in today's dollars.

Another way to see leasing's lower cost is to look at the difference in cash flows between leasing and buying. Figure 11.21 presents the timelines for both options and their difference. The top timeline details the cash flows for buying the truck, which consist of the capital expenditure today, the depreciation tax shields in the future, and the residual value. We use the present value of the residual value because it has a different discount rate than the other cash flows. The middle timeline details the lease cash flows, which consist only of the after-tax lease payments. The bottom timeline is the difference in the lease and buy cash flows (lease minus buy). What this last timeline reveals is that leasing is equivalent to a loan in which we receive \$68,765.26 today in exchange for future payments of \$17,727.70 over the next four years and a final payment of \$2,520.00 in year five.

The value of this loan at the after-tax debt cost of capital is

$$68,765.26 - \frac{17,727.70}{0.0474} (1 - (1 + 0.0474)^{-4}) - \frac{2,520.00}{(1 + 0.0474)^5} = \$3,523.09,$$

Years	0	1	2	3	4	5
CapEx	-120					
Dep. tax shield		2.52	2.52	2.52	2.52	2.52
PV(Residual value)	36.03					
Buy cash flows	-83.97	2.52	2.52	2.52	2.52	2.52

Years	0	1	2	3	4	5
After-tax lease pmts.	-15,21	-15,21	-15,21	-15,21	-15,21	
Lease cash flows	-15.21	-15.21	-15.21	-15.21	-15.21	

Years	0	1	2	3	4	5
Lease - Buy cash flows	68.77	-17.73	-17.73	-17.73	-17.73	-2.52

Figure 11.21: Lease Minus Buy Cash Flows (\$000s)

the same value we calculated above. Alternatively, we can compare the loan interest rate to the after-tax cost of debt by computing the internal rate of return (IRR) of the cash flows. Remember, the IRR is the one discount rate that sets the net present value of the cash flows equal to zero.

$$68,765.26 - \frac{17,727.70}{IRR} (1 - (1 + IRR)^{-4}) - \frac{2,520.00}{(1 + IRR)^5} = 0 \implies IRR = 2.59\%$$

The lease implied loan interest rate of 2.59% is significantly more attractive than that of the after-tax debt cost of capital of 4.74%. Finally, we can compute the present value of the loan repayments.

$$\frac{17,727.70}{IRR} (1 - (1 + IRR)^{-4}) - \frac{2,520.00}{(1 + IRR)^5} = \$65,242.17$$

This result tells us that at our debt cost of capital, the loan repayments would only allow us to borrow \$65,242.17, which is less than what the lease enables us to borrow (\$68,765.26).

Financial Lease

Now let's pretend that our lease is classified as a financial lease instead of an operating lease. In this case, the lessee gets the depreciation tax shield and can deduct the interest component of the lease payments. In other words, a financial lease is equivalent to buying and financing with an amortizing loan that starts today. Thus, the question is: Which loan is better - the one offered by our creditors or the lessor?

If we buy the asset we would borrow \$120,000 at 6% in our truck example. We would also receive the residual value in five years, which has a present value at the operating cost

of capital equal to $60,000/(1 + 0.12)^5 = \$34,045.61$. So, in effect, we would be borrowing $120,000 - 34,045.61 = \$85,954.39$.

Let's compare this figure to how much could we borrow at our debt cost of capital if the loan repayments consist of the lease payments.

$$19,250 + \frac{19,250}{1 + 0.06} + \frac{19,250}{(1 + 0.06)^2} + \frac{19,250}{(1 + 0.06)^3} + \frac{19,250}{(1 + 0.06)^4} = \$85,954.39$$

The exact same. In other words, buying by borrowing or entering into a financial lease are equivalent in this example.

Summary

To summarize, the lease vs. buy decision is fundamentally no different from other financial decisions. We want to choose the one with the highest value or lowest cost. What's tricky about leases is getting the discount rates correct. Specifically, for both the lease and depreciation tax shields, the after-tax cost of debt is the appropriate discount rate assuming the lessee expects to use depreciation tax shield each year. If the lessee is unprofitable, or there is uncertainty regarding the tax rate or the lessee's ability to utilize the depreciation tax shield that a different tax rate may be more appropriate. Likewise, the residual value at the end of the lease is typically going to be riskier than the cash flow from secured debt, in which case something other than the after-tax cost of debt, such as the after-tax WACC, may be more appropriate. In computing the WACC, we must also recognize that the project leverage is 100%.

11.8 Key Ideas

This chapter tackled the challenge of how firms (and by extension people) should finance activity. The key take-aways are the following.

1. The market value balance sheet, like a timeline, is a useful visualization to keep track of the relations between different values and expected returns/cost of capitals - cash, debt, equity, enterprise, assets.
2. When capital markets are perfect, the M&M results imply that capital structure and payout policy are both irrelevant for firm value and shareholder wealth. However, they can impact the risk of debt and equity claims, which gets reflected in expected returns. Higher leverage implies greater risk for shareholders, and to a lesser degree creditors, and therefore higher expected returns, all else equal.
3. Market imperfections, such as transaction costs, taxes, and bankruptcy costs, all provide a role for capital structure and payout policy to impact firm value and shareholder wealth. It is critical to justify any financial policy as addressing one of these imperfections in order to justify the value of financial policy.

4. Don't be fooled by popular performance metrics, such as earnings per share and price-to-earnings ratio. They are a means to an end, namely, value. Depending on the project NPV and financial strategy, EPS and P/E ratios can go up or down independently of any cost or benefit to shareholders because both value and risk are changing.
5. Existing shareholders capture the NPV of new investments when capital markets are perfect.
6. When valuing a project, the incremental financing is what determines the appropriate project cost of capital and the most convenient valuation technique. When targeting a fixed leverage ratio, WACC valuation is easiest. When the project debt policy is pre-determined, APV is easiest.

11.9 Technical Appendix

To ease the presentation, let's use the first letter of each component of the market value balance sheet: $E = Equity$, $D = Debt$, $C = Cash$, $O = Operating$, and $V = Value$.

11.9.1 Project WACCs

To derive equation 11.13 note that

$$r^E = r^O + \frac{D - C}{E} (r^O - r^D)$$

Plugging this expression into equation 11.12 and simplifying yields an expression for the WACC in terms of the operating cost of capital.

$$\begin{aligned} r^{WACC} &= \frac{D - C}{E + D - C} r^D (1 - TaxRate) + \frac{E}{E + D - C} \left(r^O + \frac{D - C}{E} (r^O - r^D) \right) \\ &= \frac{D - C}{E + D - C} r^D (1 - TaxRate) + \frac{E}{E + D - C} r^O + \frac{D - C}{E + D - C} (r^O - r^D) \\ &= r^O - \frac{D - C}{E + D - C} r^D TaxRate \end{aligned}$$

11.9.2 Cost of Capitals with Predetermined Debt Policy

When the firm follows a predetermined debt policy, the formulae for the (1) operating or unlevered cost of capital, r^O , (2) equity or levered cost of capital, r^E , and (3) weighted average cost of capital are given by the following expressions.

$$\begin{aligned} r^O &= \frac{D - T - C}{E + D - C} \times r^D + \frac{E}{E + D - C} \times r^E \\ r^E &= r^O + \frac{D - T - C}{E} \times (r^O - r^D) \\ r^{WACC} &= r^O - \frac{T}{E + D - C} (r^O - r^D) - \frac{D}{E + D - C} r^D \tau \end{aligned}$$

What's new is the variable, T , which represents the present value of the interest tax shields discounted by r^D . Because the debt is predetermined, we can compute all the future interest and interest tax shields. Discounting them by r^D and summing produces T .

For more details, see the following references.

- DeMarzo, Peter, 2007, Discounting tax shields and the unlevered cost of capital, Working paper.
- Hamada, R., 1972, "The effect of a firm's capital structure on the systematic risks of common stocks," *Journal of Finance* 27, 435-452.
- Miles, J. and J. Ezzell, 1980, "The weighted average cost of capital, perfect capital markets and project life: A clarification," *Journal of Financial and Quantitative Analysis* 15, 719-730.
- Modigliani, Franco and Merton Miller, 1958, "The cost of capital, corporation investment, and the theory of investment," *American Economic Review* 48, 261-297.
- Roberts, Michael, 2018, "Discount rates and corporate valuation," Working paper.
- Stanton, Robert and Mark Seasholes, 2005, "The assumptions and math behind WACC and APV calculations," Working paper.

11.9.3 The Effect of Default

To determine the interest rate on the loan when default is a possibility, we have two equations in need of solving. First, the expected return on the debt should satisfy the CAPM.

$$r^D = r^F + \beta^D (r^M - r^F)$$

Second, the expected return should equal the probability weighted average realized return.

$$\begin{aligned}
 r^D = & \frac{1}{3} \frac{\overbrace{[Interest + \min(Price_{Low} - Interest, Principal)]}^{\text{Cash flow to bank in Low state}}}{Principal} \\
 & + \frac{1}{3} \frac{\overbrace{[Interest + \min(Price_{Medium} - Interest, Principal)]}^{\text{Cash flow to bank in Medium state}}}{Principal} \\
 & + \frac{1}{3} \frac{\overbrace{[Interest + \min(Price_{High} - Interest, Principal)]}^{\text{Cash flow to bank in High state}}}{Principal} - 1
 \end{aligned}$$

Though this looks scary, all it says is that for each possible sales outcome, the return on debt is equal to one of two numbers depending on whether the sales price of the home is less than or greater than the amount owed to the bank. If the sales price of the home is larger than

the interest and principal, then the return to debt is just equal to the interest plus principal divided by the principal minus one. If the sales price of the home is less than the interest and principal, then the return to debt is equal to the sales price of the home divided by the principal minus one.

The loan interest rate is implicit in these two equations. It shows up in the first equation in β^D , which is the ratio of the covariance of debt returns and the market to the variance of the market. (We use the returns to the house as a proxy for the market.) The returns to debt are in turn a function of the interest rate on the loan, which determines the interest in the second equation. We can let our spreadsheet program find the interest rate such that these two equalities hold. When we do, we get a loan interest rate of 18.58%, a debt beta of 0.72, and a debt cost of capital of 6.19%.

11.9.4 Personal Taxes and Capital Structure

Consider Table 17, which shows two scenarios. The scenario on the left shows a company with \$100 of EBIT and debt financing. If it pays its all of its operating earnings to its creditors, it pays no corporate taxes and its (debt) investors receive the full \$100. However, interest is taxed as ordinary income, which means the debt investors receiving this money will have to pay income taxes on it. If the investors are wealthy individuals, they could pay as much as 37% in taxes, leaving them with only $100 \times (1 - 0.37) = \63 after taxes.

The scenario on the right side of Table Table 17 shows the same firm but with no debt. To distribute its earnings to its shareholders, the company must first pay corporate taxes, which as of 2024 were 21% for most profitable companies. This leaves the company $100 \times (1 - 0.21) = \79 to give to its shareholders. Shareholders can receive this money either through capital gains or dividends. As of 2024, long-term capital gains and qualified dividends were both taxed at 20%. So, shareholders receive $79 \times (1 - 0.20) = \$63.20$ after-tax.

Company with Debt		Company with No Debt	
EBIT	100.00	EBIT	100.00
Interest	100.00	Corporate tax	21.00
		Earnings	79.00
Debt investor		Equity investor	
Interest income	100.00	Dividend/Capital gain	79.00
Personal tax	37.00	Personal tax	15.80
After-tax income	63.00	After-tax income	63.20

Table 17: Equity Cash Flows and Returns with 95% Borrower Leverage (\$1,000s except returns)

In this example, there is a tax-disadvantage to using debt once we account for personal taxes. That said, there are a lot of assumptions behind these results. First, we assumed the corporation was paying 21% corporate tax. That need not be the case even for very profitable

companies who can rely on non-debt tax shields to reduce their taxes (not repatriating profits from low tax countries to high tax countries, receiving investment tax credits, etc.). Second, dividend and capital gains tax rates vary depending on a variety of factors including the income of the shareholder and the length of time they have been a shareholder. Third, not all investors face the same taxes. For example, pensions don't pay taxes on investment earnings whereas high income individuals may pay significant taxes. Thus, the tax benefit or disadvantage of debt financing depends on the tax rates the corporation and its investors face.

11.10 Problems

11.1 Barnes and Noble Education (BNED) had a share price of \$0.54 and 94.222 million shares outstanding as of May 21, 2024. It's most recent audited balance sheet as of April 29, 2023 is below.

(In thousands, except per share data)	April 29, 2023
Cash and cash equivalents	14,219
Receivables, net	92,512
Inventories	353,328
Prepaid expenses and other current assets	76,942
Total current assets	537,001
Property and equipment, net	68,153
Operating lease right-of-use assets	246,972
Intangible assets, net	110,632
Deferred tax assets, net	132
Other noncurrent assets	17,889
Total assets	980,779
Accounts payable	267,923
Accrued liabilities	85,759
Current operating lease liabilities	99,980
Short-term borrowings	0
Liabilities held for sale	8,423
Total current liabilities	462,085
Long-term deferred taxes, net	1,970
Long-term operating lease liabilities	184,754
Other long-term liabilities	19,068
Long-term borrowings	182,151
Total liabilities	850,028
Preferred stock	0
Common stock	551
Additional paid-in capital	745,932
Accumulated deficit	(593,356)
Treasury stock, at cost	(22,376)
Total stockholders' equity	130,751
Total liabilities and stockholders' equity	980,779

Using this information, construct a market value balance sheet for BNED and answer the following questions.

- a. What is the value of equity?
- b. What is the firm value?
- c. What is the net debt?
- d. What is the enterprise value?
- e. What is the leverage ratio? Net leverage ratio?
- f. How would you describe BNED's capital structure?

11.2 Lululemon (LULU) had a share price of \$322.98 and 117.834 million shares outstanding as of May 21, 2024. Its most recent audited balance sheet as of January 28, 2024 is below.

(Amounts in millions)	<u>January 28, 2024</u>
Cash and cash equivalents	2,244
Accounts receivable, net	125
Inventories	1,324
Prepaid and receivable income taxes	184
Prepaid expenses and other current assets	185
Total current assets	<u>4,061</u>
Property and equipment, net	1,546
Right-of-use lease assets	1,266
Goodwill	24
Intangible assets, net	0
Deferred income tax assets	9
Other non-current assets	187
Total assets	<u>7,092</u>
Accounts payable	348
Accrued liabilities and other	349
Accrued compensation and related expenses	326
Current lease liabilities	249
Current income taxes payable	12
Unredeemed gift card liability	306
Other current liabilities	40
Total current liabilities	<u>1,631</u>
Non-current lease liabilities	1,154
Non-current income taxes payable	16
Deferred income tax liabilities	30
Other non-current liabilities	29
Total liabilities	<u>2,860</u>
Preferred stock	0
Common stock	1
Additional paid-in capital	575
Retained earnings	3,920
Accumulated other comprehensive loss	(264)
Total shareholders equity	<u>4,232</u>
Total liabilities and shareholders equity	<u>7,092</u>

Using this information, construct a market value balance sheet for LULU and answer the following questions.

- a. What is the value of equity?
- b. What is the firm value?
- c. What is the net debt?
- d. What is the enterprise value?
- e. What is the leverage ratio? Net leverage ratio?

f. How would you describe LULU's capital structure?

- 11.3 Apple (AAPL) had a share price of \$191.04 and 15.442 billion shares outstanding as of May 21, 2024. Its equity beta was 1.26, its debt cost of capital was 4.2%, and its most recent audited balance sheet as of September 30, 2023 is below. The yield on a 10-year Treasury note was 4.45% and the market risk premium was 5.00%.

(Amounts in millions)	<u>September 30, 2023</u>
Cash and cash equivalents	29,965
Marketable securities	31,590
Accounts receivable, net	29,508
Vendor non-trade receivables	31,477
Inventories	6,331
Other current assets	14,695
Total current assets	<u>143,566</u>
Marketable securities	100,544
Property, plant and equipment, net	43,715
Other non-current assets	64,758
Total assets	<u>352,583</u>
Accounts payable	62,611
Other current liabilities	58,829
Deferred revenue	8,061
Commercial paper	5,985
Term debt	9,822
Total current liabilities	<u>145,308</u>
Term debt	95,281
Other non-current liabilities	49,848
Total liabilities	<u>290,437</u>
Common stock	73,812
Accumulated deficit	(214)
Accumulated other comprehensive loss	(11,452)
Total shareholders' equity	<u>62,146</u>
Total liabilities and shareholders' equity	<u>352,583</u>

Using this information, construct a market value balance sheet for AAPL and answer the following questions.

- What is the firm value?
- What is the enterprise value?
- What is the equity cost of capital according to the CAPM?
- What is the operating cost of capital? If this is different from the equity cost of capital, why is one higher than the other?
- What is the asset cost of capital? If the asset cost of capital is different from the operating cost of capital, why is one higher than the other?

- f. Can you estimate the equity cost of capital as a function of the asset cost of capital? Operating cost of capital? How do your answers compare to the estimate from question c.?
- g. If Apple wants to assess a new project whose cash flows exhibit risks similar to the rest of the business, which discount rate should it use assuming capital markets are perfect?
- h. What are the betas of debt, cash, operating assets, and total assets? How does each compare to the betas of the market and the risk-free asset?
- 11.4 Continuing the previous problem, assume that Apple's (AAPL) projected 2024 operating income (EBIT) and tax rate were \$114.301 billion and 14.7%. Answer the following questions assuming taxes are the only market imperfection.
- a. How much debt should Apple issue at 5.8% to completely eliminate any taxes? Assume that *all* interest is tax deductible and that any new debt can be issued at 5.8%. How large is the interest tax shield?
- b. Answer the previous questions assuming that Apple must abide by current tax law and can only deduct interest up to 30% of its operating income.
- 11.5 Microsoft is currently trading at \$429.04 per share with 7.44 billion shares outstanding. Its equity beta is 0.89 and its debt cost of capital is 4.3%. The current yield on a 10-year Treasury note is 4.45% and the market risk premium is 5.00%. Microsoft's market value balance sheet looks as follows.

Assets	Claims
Cash	94.8 Debt 79.9
Operating	3,177.1 Equity 3,192.1
Total assets	3,272.0 Total claims 3,272.0

Activist investors have been pressuring Microsoft's CFO, Amy Hood, to reduce the company's surplus of cash. In response, Amy has decided to investigate how distributing cash would affect shareholders.

Using this information, answer the following questions.

- a. What is Microsoft's current equity cost of capital before any distribution?
- b. What is Microsoft's current after-tax WACC before any distribution and assuming Microsoft maintains a target capital structure.
- c. If Amy distributes all of the cash to shareholders via a one-time special dividend, what are the implications for Microsoft shareholder value and expected returns in a perfect capital market setting? Answer this question by reconstructing the market value balance sheet after the dividend and estimating the equity cost of capital. What are the implications for the after-tax WACC assuming taxes are the only imperfection and Microsoft intends to maintain this new capital structure?

- d. If Amy instead retires all of the debt, keeping any residual cash on the balance sheet, what are the implications for Microsoft shareholder value and expected returns? Answer this question by reconstructing the market value balance sheet after the debt retirement and estimating the equity cost of capital. What are the implications for the after-tax WACC assuming taxes are the only imperfection and Microsoft intends to maintain this new capital structure?
- e. Considering only taxes, what should Amy do with the cash on Microsoft's balance sheet and why?

11.6 Bear Industries is a construction firm specializing in municipal infrastructure. Bear is in need of a concrete grinder, which costs \$100,000 and has an estimated life of eight years. They are considering a one-year lease of the grinder at the end of which the grinder will have a residual value of \$80,000. Bear has an annual debt cost of capital equal to 8% and an annual operating cost of capital equal to 14%.

Using this information, answer the following questions assuming perfect capital markets.

- a. What should the monthly lease payments be assuming payments are made at the start of each month and the residual value has a risk profile similar to that of Bear's debt?
- b. The CFO has argued that the residual value is, in fact, substantially riskier than Bear's debt because of the cyclical nature of government spending. What should the the monthly lease payment be under this different risk profile? How does it compare to you answer to the previous question and explain any difference?
- c. What is the monthly loan payment if instead Bear decides to purchase the grinder by borrowing money for one year? Assume the loan payments are made at the *end* of each month. How do the loan payments compare to the lease payments? Explain any difference.

11.7 Terren General Contractors is in need of a backhoe loader, which costs \$150,000 and has a depreciable life of 10 years after which the backhoe will be worthless. Terren is deciding between purchasing and entering into a sale-type lease with Caterpillar, the manufacturer of the backhoe. If they choose to lease the backhoe, the lease term is 10 years with annual payments of \$21,500, and Terren would be responsible for all maintenance, upkeep, and insurance on the backhoe (i.e., net lease). Terren's debt cost of capital is 8%, and their tax rate is 21%.

Using this information, answer the following questions.

- a. Assuming the backhoe is depreciated on a straight-line basis, what is the annual depreciation?
- b. What are the incremental cash flows for buying and leasing the backhoe?
- c. What interest rate is implicit in the lease-equivalent loan? Is it high or low and relative to which benchmark?

- d. How much value does leasing create or destroy relative to buying? Should Terren buy or lease the backhoe?
- e. What is the annual lease payment that would make Terren indifferent between buying and leasing?
- f. Terren has asked Caterpillar to shorten the lease to two years with annual payments of \$80,000. How much value does leasing create or destroy relative to buying under these new terms? How do you think the tax authority would view this contract?

11.8 Coral Minerals is a mining company registered in the Cayman islands where there is no corporate tax. Coral needs a large excavator for a silicon mine. The excavator costs \$650,000 and has a 12 year depreciable life and no residual value at the end of its life. Coral's debt cost of capital is 10% and is considering a 12 year lease with annual payments of \$86,000.

Using this information, answer the following questions.

- a. Should coral buy or lease the excavator? Justify your answer by estimating the NPV and IRR of the lease equivalent loan.
- b. Does it matter whether the lease is classified as an operating or financial lease? Explain.

Chapter 12

Corporate Valuation (Incomplete)

Fundamental value relation

$$Value_t = \frac{CashFlow_{t+1}}{(1+r)} + \frac{CashFlow_{t+2}}{(1+r)^2} + \frac{CashFlow_{t+3}}{(1+r)^3} + \dots$$

Corporate valuation is just that - the valuation of corporations or business entities, more generally. It's the perfect topic to end the book for two reasons. First, it encompasses lessons from all of the previous chapters and, as such, adds a capstone project to the text. Second, corporate valuation is a central component of many applications in finance include the following.

- Mergers and acquisitions (M&A)
- Leveraged buyouts (LBOs)
- Financial planning and analysis (FP&A)
- Raising capital (IPOs, SEOs, Debt offerings, Loan originations)
- Debt and equity analysis
- Due diligence in fairness opinions
- Determination of recovery in bankruptcy
- Fundamental investing

- Buy- and sell-side analysis
- Partnership dissolution
- Restructurings

Our illustrative vehicle for this chapter is a hypothetical private dental practice - VB Dental. The practice has been owned and operated for 20 years by, Bill Lemanski, D.D.S. As of the end of 2021, Bill was considering selling the business and wanted to know its value. The objective of this exercise is to estimate the value both from Bill's perspective as well as that of a private equity investor. To do so, we'll employ different valuation techniques, including discounted cash flow analysis (DCF) and **market multiples** or the **method of comparables** and explore the effects of different financial strategies.

12.1 Business Model

Bill's business model is relatively straightforward. He generates revenue by seeing adult - over 18 years of age - patients for dental services. Providing these services cost him money in terms of expenses and investment. His most recent income statement and balance sheet for 2021 are presented in Figure 12.1. The business earned a little over \$1.1 million off \$2.0 million in sales. His assets have a book value of \$3.2 million, which was finance with \$2.7 million of debt.

P&L (\$000)	2021	Balance sheets (\$000s)	2021
Sales	2,000.0	Cash	438.2
Expenses	584.2	Accounts receivable	500.0
EBITDA	1,415.8	Current assets	938.2
Depreciation	185.0	Net PP&E	2,355.0
EBIT	1,230.8	Total assets	3,293.2
Interest	99.0		
EBT	1,131.8	Accrued wages	17.0
Taxes	0.0	Current liabilities	17.0
Earnings	1,131.8	Long-term debt	2,668.4
		Total liabilities	2,685.4
		Shareholders equity	607.7
		Total liabilities & shareholder equity	3,293.2

Figure 12.1: Historical P&L Statement and Balance Sheet for VB Dental

Because Bill is both the owner and only shareholder, he has structured his dental practice as a **sole proprietorship** or **S-corporation** (S-corp). The choice of business structure has

important legal and tax implications for the owners and managers of the business. There are many different types in the U.S. from which to choose, though with different limitations and implications. Some of the more common structures include the following.¹

- **Partnerships.** Partnerships are business owned by two or more individuals. **General partnerships (GP)** are where partners share in both the management of the business and its profits, and have unlimited liability for business debts. In other words, creditors can come after both the corporate and personal assets of the partners to collect debts. Some medical practices and architectural firms are structured as general partnerships. **Limited partnerships (LP)** consist of both general and limited partners. The general partners (GPs) manage the business and have unlimited liability. Limited partners do not participate in the management of the business and have limited liability. Most private equity companies are organized as limited partnerships. **Limited liability partnerships (LLP)** are similar to a general partnership in that the partners all share in the management and profits of the business. However, all of the partners are protected by limited liability. That is, creditors can only go after the assets of the business, not the personal assets of the partners. Many law firms are organized as LLPs.
- **Corporation.** Corporations are distinct from their owners in the eyes of the law. (**C-corporations** or **C-corps**) can raise capital by issuing stock and its owners, i.e., shareholders, have limited liability. However, shareholders are subject to **double taxation** in that the corporation pays corporate taxes and then the owners pay taxes on any profits they receive through dividends or capital gains. Most publicly traded companies and a number of very large privately held companies are organized as C-corporations. **S-corporations (S-corp)** are similar to C-corps with two key differences. First, the number of shareholders is limited to 100. Second, S-corps pay no federal taxes and, as such, avoid double taxation at least at the federal level. Many smaller, privately owned businesses are structured as S-corps.
- **Limited liability company (LLC).** An LLC is a mix of the previous two classes of business structure. Owners or **members** have limited liability and don't pay taxes on profits at the federal level. LLCs also have a flexible management structure. Many real estate investors and consulting businesses are structured as LLCs.

¹Other business structures include nonprofit corporations, cooperatives, professional corporations, joint ventures, and benefit corporations (B-corporations).

The choice of legal structure is an important one for companies, but beyond the scope of this book.

12.2 Discounted Cash Flow (DCF) Analysis

12.2.1 Valuation Strategy

Our DCF valuation strategy is illustrated in Figure 12.2. We're going to take a multi-stage approach in which he first constructs detailed pro forma P&L statements, balance sheets, and cash flow forecasts over the next five years. The dynamics and growth of his business over this period are entirely flexible and intended to represent, as accurately as possible, the evolution of his business during this time. We've chosen a five-year horizon for two reasons. First, five years from now his business growth is expected to reach a sustainable or **steady-state** rate of 4% per year. Second, the five year horizon provides Bill with a long enough perspective to make consequential decisions such as investment or possible retirement.

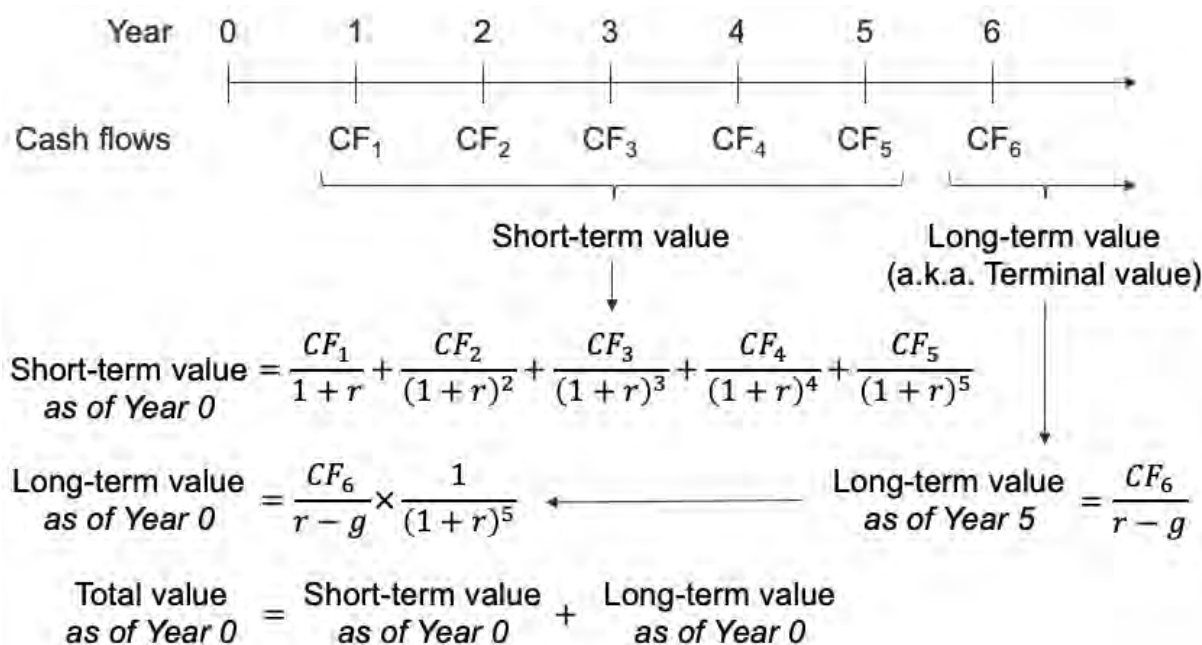


Figure 12.2: Discounted Cash Flow Analysis Strategy

By discounting the projected cash flows and adding, we can obtain the **short-term value** of his business today. This value represents the present value of the next five years of business operations. Mathematically, short-term value as of today can be expressed like so.

$$\text{Short-term value in year 0} = \frac{CF_1}{1+r} + \frac{CF_2}{(1+r)^2} + \frac{CF_3}{(1+r)^3} + \frac{CF_4}{(1+r)^4} + \frac{CF_5}{(1+r)^5} \quad (12.1)$$

The short-term value is *not* the same as the value of his business. It's just part of that value. As suggested by Figure 12.2, the other part is the long-term value generated by cash flows from year six onward. Because forecasting an infinite number of cash flows is infeasible, we'll forecast the year six cash flow and then assume all subsequent cash flows grow at a constant sustainable or **steady-state** rate of 4% per year. Coupled with the assumption of a constant discount rate, r , the **long-term value** of his business as of the **terminal year** - year five in this example - can be computed as the value of a growing perpetuity (equation ??).

$$\text{Long-term value in year 5} = \frac{CF_6}{r - g} \quad (12.2)$$

The long-term value is more commonly referred to as **terminal value** or **continuation value**. It represents our estimate of the value of Bill's business five years from today. So, if Bill was planning on selling his business in five years and wanted to know its value at that future date, the terminal value provides an estimate.

Because the terminal value is the value as of five years from today, we can't add it to the short-term value as of today. The two estimates have different time units. We have to discount the terminal value back to today.

$$\text{Long-term value in year 0} = \frac{CF_6}{r - g} \times \frac{1}{(1 + r)^5} \quad (12.3)$$

Now we can add the results of equations 12.1 and 12.3 to get the total value of Bill's business today. Let's summarize this process.

DCF Recipe for Corporate Valuation

1. Forecast cash flows over T periods - the short-term - where the number of periods T is determined by the length of time required to achieve a sustainable growth rate in perpetuity. Over this horizon, we want to forecast all of the cash flow value drivers - sales, expenses, investment, taxes, and financial policy - in great detail to ensure we understand the business and its future evolution. This step is important for budgeting, planning, and financing.
2. Discount each cash flow over the short-term and sum to obtain the short-term value of the business, which reflects the value generated by operations over the next T periods.

$$\text{Short-term value in period 0} = \frac{CF_1}{1 + r} + \frac{CF_2}{(1 + r)^2} + \dots + \frac{CF_T}{(1 + r)^T} \quad (12.4)$$

- Forecast the cash flow one period after the terminal year, CF_{T+1} , and use the growing perpetuity formula (equation ??) to compute the value of the business as of the terminal period T .

$$\text{Long-term value in period } T = \frac{CF_{T+1}}{r - g} \quad (12.5)$$

This is the terminal or continuation value of the business.

- Discount the terminal value back to today, year 0, to get the present value of the terminal value.

$$\text{Long-term value in period 0} = \frac{CF_{T+1}}{r - g} \times \frac{1}{(1 + r)^T} \quad (12.6)$$

- Add the short-term value as of today (equation 12.4) to the long-term value as of today (equation 12.6) to get the total value of the business as of today.

12.2.2 Short-Term Value

To estimate the short-term value, we need to forecast the free cash flows. Recall the definitions of unlevered (FCF) and levered, or equity (FCFE), free cash flows.

$$FCF = (Sales - Expenses - D\&A) \times (1 - TaxRate) + D\&A - NLTI + NWCI \quad (12.7)$$

$$FCFE = FCF - (1 - TaxRate) \times Interest + NDI \quad (12.8)$$

$S\&A$ represents depreciation and amortization, and any other non-cash expenses such as depletion. $NLTI$ is net long-term investment and includes capital expenditures, intangible investments, acquisitions, and asset sales. $NWCI$ is net working capital investment or the change in net working capital. $Interest$ is the interest expense from any debt. NDI is net debt issuance or debt issuances less any debt retirements. The relevant $TaxRate$ in this setting is the effective tax rate because the earnings correspond to the entire earnings of the company.²

Forecast Drivers

To estimate the free cash flows, we need to detail the forecast drivers - the assumptions governing the evolution of each component of free cash flow.

²Recall from chapter 5 that when we are assessing capital projects we use the marginal tax rate for $TaxRate$ because the earnings are incremental to the business.

- **Sales.** Bill's revenue comes from two streams: care and cosmetics. Dental care consists of preventative treatments such as regular cleanings and check-ups, fluoride treatments, dental sealants, and oral cancer screenings. It also includes restorative care such as filling cavities, installing crowns and bridges, performing root canals, and creating dentures. Revenue from care services are expected to grow at 5% per year. Cosmetic dental procedures include teeth whitening and bonding, veneer placement, and minor teeth straightening procedures. Revenue from cosmetic services are expected to grow at 10% per year as demand for these services increases.
- **Expenses.**
 1. *Receptionists.* There are two receptionists each costing \$78,000 in 2021. These costs are expected to grow by 3% per year.
 2. *Hygienists.* Bill currently has two hygienists - one junior and one senior. In 2021, the former cost \$110,500; the latter cost \$141,700. Three years from now in 2024, Bill plans to hire a third hygienist at \$110,500 for the year to lighten his work load. All three hygienist costs are expected to increase by 5% year because of the growing demand for their services.
 3. *Rent.* Rent for medical office space in 2021 was \$125,000. Bill has a long-term operating lease with the building owner limiting his annual rental increases to 1.5%.
 4. *Insurance.* Bill maintains several insurance policies for his business including professional liability (malpractice), general liability, property for his equipment, workers compensation, and health insurance. The total annual cost for this insurance was \$15,000 in 2021 and is expected to increase by 4% per year.
 5. *Advertising.* Bill advertises his practice through several channels including print and social media. His annual advertising budget in 2021 was \$8,000, which he expects will increase by 5% annually.
 6. *Supplies.* Bill purchased \$18,000 in dental supplies in 2021. Competition in this space is limiting price growth to 3% per year.
 7. *Equipment maintenance.* Bill's 2021 equipment (X-ray machine, dental chairs, computers, etc.) maintenance expenses totaled \$14,000. He expects these costs to increase by 3% per year.
- **Net long-term investment and depreciation.** Three years from today, in 2024, Bill will need to replace aging equipment such as his X-ray machine and IT infrastructure.

He also plans on modernizing the decor. Total capital expenditures that year are estimated to be \$500,000 in total. Annual depreciation over the five year horizon is estimated at 7.4% of the start of year capital stock.

- **Working capital.**

1. *Required cash.* Bill maintains cash equal to 75% of his current total expenses.³
2. *Accounts receivable.* His days receivable is 90 days, due mostly to insurance processing.
3. *Accrued wages.* He pays his employees - receptionists and hygienists - twice a month only *after* they have worked. This wage policy results in approximately 15 days payable on his labor costs. Accrued wages are similar in spirit to accounts payable. They are a current liability except instead of receiving credit from suppliers, Bill is receiving credit from his employees.

- **Taxes.** As an S-corp in Florida, Bill's business pays no taxes at the federal or state level. The corporate tax rate is zero. (He does pay personal income taxes on the equity free cash flow he receives from the business.)

The unlevered free cash flow forecast drivers are summarized in Figure 12.3. Historical data is identified in the 2021A - "A" for actual - column. The forecasted data is in the columns 2022E through 2026E - "E" for estimated. The step column controls the year-to-year changes of each assumption. Blue font identifies numbers that are hard-coded or hand-entered. Black font identifies formulas. Despite appearances, there is no rounding in any computations. The numbers are only formatted to ease analysis of the model. Recall from chapter ?? that we **never round numbers in a financial model**. Small rounding errors can quickly become large errors in computations.

Unlevered Free Cash Flows

We now have to translate our forecast drivers into free cash flow forecasts. This starts with a pro forma P&L presented in Figure 12.4. Sales are the sum of care and cosmetic revenue streams. Expenses are the sum of all labor, rent, insurance, advertising, supplies, and maintenance costs. EBITDA is Sales less Expenses.

³A more sensible assumption is that required cash at the end of each period equal 75% of *next year's* total expenses to ensure the business has sufficient liquidity for the expenses for which the cash is intended. However, this would require an estimate of expenses after the terminal year, year five, which while not difficult adds an unnecessary complication to the exercise.

Forecast drivers (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Sales						
Care	1,250.0	1,312.5	1,378.1	1,447.0	1,519.4	1,595.4
Cosmetic	750.0	825.0	907.5	998.3	1,098.1	1,207.9
Expenses						
Salaries + benefits						
Reception 1	78.0	80.3	82.8	85.2	87.8	90.4
Reception 2	78.0	80.3	82.8	85.2	87.8	90.4
Hygenist 1	110.5	116.0	121.8	127.9	134.3	141.0
Hygenist 2	141.7	148.8	156.2	164.0	172.2	180.8
Hygenist 3				110.5	116.0	121.8
Rent	125.0	126.9	128.8	130.7	132.7	134.7
Insurance	15.0	15.6	16.2	16.9	17.5	18.2
Advertising	8.0	8.4	8.8	9.3	9.7	10.2
Cleaning and medical supplies	18.0	18.5	19.1	19.7	20.3	20.9
Equipment maintenance	10.0	10.3	10.6	10.9	11.3	11.6
Capital expenditures				500.0		
Depreciation (% start net PPE)	7.4%	7.4%	7.4%	7.4%	7.4%	7.4%
Working capital						
Days in period	360	360	360	360	360	360
Required cash (% total exp.)	75.0%	75.0%	75.0%	75.0%	75.0%	75.0%
Days receivable	90	90	90	90	90	90
Days payable (wages)	15	15	15	15	15	15
Tax rate	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Figure 12.3: Unlevered Free Cash Flow Drivers

P&L (\$000)	2021A	2022A	2023A	2024A	2025A	2026A
Sales	2,000.0	2,137.5	2,285.6	2,445.3	2,617.5	2,803.2
Expenses	584.2	605.2	627.1	760.4	789.6	820.1
EBITDA	1,415.8	1,532.3	1,658.5	1,684.9	1,827.8	1,983.1
Depreciation	185.0	174.3	161.4	149.4	175.4	162.4
EBIT	1,230.8	1,358.0	1,497.2	1,535.5	1,652.5	1,820.7
Taxes	0.0	0.0	0.0	0.0	0.0	0.0
NOPAT	1,230.8	1,358.0	1,497.2	1,535.5	1,652.5	1,820.7

Figure 12.4: Pro Forma P&L Statements

Depreciation is computed in concert with net long-term investment (NLTI) and is shown in Figure 12.5. Bill's business entered 2022 - ended 2021 - with \$2.355 million worth of plant, property, and equipment net of accumulated depreciation (net PP&E). He anticipates no new capital expenditures during 2022, which would add to the net PP&E. However, 7.4% of his start of period capital stock - \$174,300 - is expected to depreciate in 2022. At the end

of the 2022, he should have $2,355 - 174.3 = \$2,180,700$ of net PP&E. More generally,

$$\text{Ending Net PP\&E} = \text{Beginning Net PP\&E} + \text{Capital Expenditures} - \text{Depreciation}.$$

Capex/depreciation schedule (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Beginning net PP&E	2,500.0	2,355.0	2,180.7	2,019.4	2,369.9	2,194.5
Capex	40.0	0.0	0.0	500.0	0.0	0.0
Depreciation	185.0	174.3	161.4	149.4	175.4	162.4
Ending net PP&E	2,355.0	2,180.7	2,019.4	2,369.9	2,194.5	2,032.2

Figure 12.5: Capital Expenditure and Depreciation Schedule

Deducting our depreciation estimates from EBITDA produces the EBIT forecast in our pro forma P&L, Figure 12.4. Because there are no corporate taxes, net operating profit after taxes (NOPAT) is equal to EBIT.

At this point, we have all but one piece of information needed to estimate unlevered free cash flows - net working capital investment. Figure 12.6 presents the working capital schedule. Required cash is 75% of total expenses. In 2022, the ending cash balance is $0.75 \times 605,200 = \$453,900$. Assuming sales are uniformly distributed during the period (year), accounts receivable can be computed like so.

$$\text{Accounts receivable} = \frac{\text{Days receivable}}{\text{Days in period}} \times \text{Sales in period}$$

At the end of 2022, Bill's accounts receivable is estimated to be $90/360$ times $2,137,500 = \$534,375$. Accrued wages is estimated similarly to accounts payable only using wage expenses. For 2022, Bill's ending accrued wages are $15/360 \times 425,400 = \$17,725$. Net working capital is equal to current assets minus current liabilities. However, we are interested in the how the amount is *changing* from period to period, which corresponds to the investment in net working capital. Clear from the last row in Figure 12.6, Bill is increasing his investment in net working capital to keep up with his growing top line (i.e., sales).

Figure 12.7 presents the free cash flow schedule and our estimate of unlevered free cash flows using equation 12.7. These cash flows correspond to the money Bill can use to pay off his debt or distribute to himself as the sole shareholder.

Financial Policy and Cost of Capital

To discount and sum the unlevered free cash flows, we need a discount rate, r . The appropriate discount rate is the Operating, or unlevered, cost of capital, r^O . We'll show how to

Working capital schedule (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Required cash	438.2	453.9	470.3	570.3	592.2	615.1
Receivables	500.0	534.4	571.4	611.3	654.4	700.8
Current assets	938.2	988.3	1,041.7	1,181.6	1,246.6	1,315.9
Accrued wages	17.0	17.7	18.5	19.3	20.1	20.9
Current liabilities	17.0	17.7	18.5	19.3	20.1	20.9
Net working capital	921.1	970.6	1,023.2	1,162.3	1,226.5	1,295.0
Change	46.8	49.4	52.7	139.1	64.2	68.5

Figure 12.6: Working Capital Schedule

Free cash flow (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
NOPAT (EBIT x (1-t))		1,358.0	1,497.2	1,535.5	1,652.5	1,820.7
Depreciation		174.3	161.4	149.4	175.4	162.4
Capex		0.0	0.0	500.0	0.0	0.0
Change in NWC		49.4	52.7	139.1	64.2	68.5
Unlevered free cash flow		1,482.9	1,605.9	1,045.8	1,763.7	1,914.6

Figure 12.7: Free Cash Flow Schedule

estimate this rate later when we discuss market multiples. For now, we'll assume it is 9.41% in which case the short-term value of Bill's company as of today is equal to

$$\text{Short-term value in year 0} = \frac{1,428,887}{1 + 0.0941} + \frac{1,605,863}{(1 + 0.0941)^2} + \dots + \frac{1,914,627}{(1 + 0.0941)^5} = \$5,946,829$$

According to Bill's balance sheet (Figure ??), he has \$2.719 million of long-term debt. If we want to account for the interest tax shield this debt provides, we need to discount the cash flows by the weighted average cost of capital or WACC. Although there is no tax shield in this example because the corporate tax rate is zero, it's instructive to estimate the WACC and ensure it is equals the unlevered cost of capital, r^O .

From chapter 11, the WACC equation will differ depending on the firm's financial policy - how Bill plans on using debt financing. Let's assume he wants to maintain a constant target leverage ratio of 10%. Let's also assume that his debt cost of capital is 5.50%. From equation 11.13 in chapter 11, his weighted average cost of capital (WACC) can be expressed as follows.

$$r^{WACC} = r^O - \frac{Debt - Cash}{Equity + Debt - Cash} \times TaxRate \times r^D$$

However, while Bill has cash on his balance sheet, \$438,200 as of 2021, he has no *excess* cash. It is all *required* cash, part of his working capital. Because excess cash is zero, the return on

operating assets, r^O , is equal to the return on assets, r^A , and the WACC can be expressed as in equation 11.11 from chapter 11.

$$r^{WACC} = r^A - \frac{Debt}{Equity + Debt} \times TaxRate \times r^D$$

Plugging numbers in produces a WACC of

$$r^{WACC} = 0.0941 - 0.10 \times 0 \times 0.055 = 0.0941.$$

With no tax expense, there is no tax shield from using debt and the WACC and unlevered cost of capital are equal. Were the tax rate to increase from zero, Bill's interest payments would generate a tax shield what would reduce the WACC and increase the value of his firm. Though, the cost of the tax payments themselves would outweigh any benefit from this interest tax shield.

12.2.3 Long-Term Value

Our next step is to estimate the long-term, or terminal, value of the business. This is the value of the cash flows from year six onward. As we discussed above, we will estimate this value using our growing perpetuity formula, equation 12.5. This formula requires three inputs: The cash flow in year $T + 1$, the sustainable growth rate g , and the discount rate, r . Above we assumed a sustainable growth rate of 4% per year. The unlevered operating cost of capital is unchanged from 9.41% assuming no significant change to the business risk. Assuming Bill maintains a constant 10% target leverage ratio, the WACC will be unchanged as well.

This leaves estimation of the cash flow in year six, CF_6 . One approach is to apply the sustainable growth rate of 4% to the unlevered free cash flow in year five. Doing so implies an unlevered free cash flow in year six equal to $1,914,627 \times (1 + 0.04) = \$1,991,212$. However, this ignores the long-term and net working capital investment required to generate the 4% growth. To account for this required investment, we can estimate the cash flow one year after the terminal year like so.

$$CF_{T+1} = NOPAT_T \times (1 + g) - \text{Long-term assets}_T \times g - \text{Net working capital}_T \times g \quad (12.9)$$

The first term in equation 12.9 estimates the NOPAT one year after the terminal year by growing after-tax earnings at the sustainable growth rate. The second and third terms correspond to the investment needed to generate this growth in earnings. The second term

multiplies the amount of long-term assets - tangible and intangible - on the balance sheet in the terminal year by the growth rate. In our dental example, long-term assets consists of net PP&E in the fifth year. The third term multiplies the amount of net working capital in the terminal year by the growth rate.

Plugging values into equation 12.9 yields the following estimate.

$$\begin{aligned} CF_6 &= 1,820,706 \times (1 + 0.04) - 2,032,153 \times 0.04 - 1,294,960 \times 0.04 \\ &= \$1,760,450 \end{aligned}$$

This estimate is substantially less than what we obtained when applying the growth rate to the terminal year cash flow (\$1,991,200) precisely because it accounts for the investment required to generate growth.

With all of the components, we can estimate the long-term value of Bill's business five years from today in 2026 using equation 12.2.

$$\text{Long-term value in year 5} = \frac{1,760,450}{0.0941 - 0.04} = \$32,517,635$$

Discounting this value back to today, end of 2021, yields

$$\text{Long-term value in year 0} = \frac{32,517,635}{(1 + 0.0941)^5} = \$20,737,568.$$

12.2.4 Total Value

The total enterprise (i.e., operating asset) value of Bill's business as of 2021 is the sum of the short-term and long-term values as of 2021.

$$5,946,829 + 20,737,568 = \$26,684,397$$

Bill's market value balance sheet as of the end of 2021 is show in Table 1. As mentioned above, all cash on Bill's balance sheet is part of net working capital, which is a part of operating assets. So, there is no excess cash on the market value balance sheet. Thus, the total asset value of his business is equal to the operating asset value. The market value of debt is assumed to be reasonably well approximated by the book value because interest rates have not changed significantly since Bill borrowed the money and his business is in good financial health. Bill's equity value is obtained by subtracting the debt value from the total asset value.

Assets		Claims	
Cash, C	0.0	Debt, D	2,668.4
Operating, O	26,684.4	Equity, E	24,016.0
Total assets, V	26,684.4	Firm claims, V	26,684.4

Table 1: Market Value Balance Sheet - VB Dental 2021. Values in \$000s

We've estimated the value of Bill's business today (2021) and five years from today in the terminal year (2026). But, we can estimate the value of his business at any time with the Fundamental Value Relation or with our recursive result, equation 11.14 from chapter 11. For example, the value of Bill's business four years from today in 2025 can be computed like so.

$$Value_4 = \frac{CashFlow_5 + Value_5}{1 + r} = \frac{1,914,627 + 32,517,635}{1 + 0.0941} = \$31,469,753$$

The value three years from today in 2024 is

$$Value_3 = \frac{CashFlow_4 + Value_4}{1 + r} = \frac{1,763,683 + 31,469,753}{1 + 0.0941} = \$30,374,072.$$

And so on. Figure 12.8 presents all of the estimated enterprise values from 2021 through 2026.

Valuation (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Enterprise value	26,684.4	27,713.5	28,716.6	30,374.1	31,469.8	32,517.6

Figure 12.8: Enterprise Values (\$000s)

12.2.5 Debt Policy and Equity Values

With the estimated enterprise values, we can compute Bill's debt policy. Specifically, we can estimate how much debt he must have outstanding each year to ensure he hits his 10% target leverage ratio. These debt balances will allow us to estimate how much debt he must issue or repurchase each year. Figure 12.9 presents these estimates along with the enterprise values.

To understand the schedule, let's focus on 2022. Bill begins the year with \$2.668 million dollars of debt. At a 5.5% debt cost of capital, this creates an interest expense of $2,668,440 \times 0.055 = \$146,764$. At the end of 2022, the enterprise value of the business is \$27,713,534. To maintain a target leverage ratio of 10%, Bill needs $0.10 \times 27,713,534 = \$2,771,353$ of

Valuation (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Enterprise value	26,684.4	27,713.5	28,716.6	30,374.1	31,469.8	32,517.6
Debt schedule	2021A	2022E	2023E	2024E	2025E	2026E
Beginning debt	1,800.0	2,668.4	2,771.4	2,871.7	3,037.4	3,147.0
Net debt issue	868.4	102.9	100.3	165.7	109.6	104.8
Ending debt	2,668.4	2,771.4	2,871.7	3,037.4	3,147.0	3,251.8
Interest expense	99.0	146.8	152.4	157.9	167.1	173.1

Figure 12.9: Debt Schedule (\$000s)

debt outstanding at the end of the year. Therefore, he must issue or borrow an additional $2,771,353 - 2,668,440 = \$102,914$ of debt in 2022. Subsequent year calculations are similar. As his enterprise value changes, the amount of debt he must hold to maintain his target leverage ratio changes, meaning Bill must issue or retire debt each year.

With both enterprise and debt values, equity values follow from the balancing feature of the market value balance sheet.

$$\text{Equity} = \text{Enterprise Value} - \text{Debt} + \text{Cash} \quad (12.10)$$

Because Bill doesn't hold any excess cash, Cash equals zero every year. The year-on-year market value balance sheets are presented in Figure 12.10. For comparison, the year-on-year book balance sheets are presented as well. It's interesting to note the relatively large amount of debt financing when compared to the book value of equity. We'll explore this more deeply in the next section.

With the debt values we can also construct a full P&L statement, complete with interest expense and taxes that are based on pre-tax income (EBT), as opposed to EBIT. Figure 12.11 presents the full P&L pro forma income statements for Bill's business.

Finally, the debt policy also allows us to compute the equity or levered free cash flows, i.e., the money available to or need from Bill. Recall our definition of equity free cash flow, $FCFE$, from equation 12.8. We start with unlevered free cash flow and then deduct the after-tax interest expense and add any debt issuances net of retirements. Figure 12.12 presents the levered free cash flows. These figures represent Bill's annual gross income. (He has to pay personal income taxes on this money.)

Notice that Figure 12.12 also presents the present values of the levered cash flow and their sum or short-term equity value (\$5,717,665). We estimated Bill's equity value above

Balance sheets (\$000s)	2021A	2022E	2023E	2024E	2025E	2026E
Market value balance sheet						
Operating assets	26,684.4	27,713.5	28,716.6	30,374.1	31,469.8	32,517.6
Total assets	26,684.4	27,713.5	28,716.6	30,374.1	31,469.8	32,517.6
Debt	2,668.4	2,771.4	2,871.7	3,037.4	3,147.0	3,251.8
Equity	24,016.0	24,942.2	25,844.9	27,336.7	28,322.8	29,265.9
Total claims	26,684.4	27,713.5	28,716.6	30,374.1	31,469.8	32,517.6
Book balance sheet						
Cash	438.2	453.9	470.3	570.3	592.2	615.1
Accounts receivable	500.0	534.4	571.4	611.3	654.4	700.8
Current assets	938.2	988.3	1,041.7	1,181.6	1,246.6	1,315.9
Net PP&E	2,355.0	2,180.7	2,019.4	2,369.9	2,194.5	2,032.2
Total assets	3,293.2	3,169.0	3,061.1	3,551.5	3,441.1	3,348.1
Accrued wages	17.0	17.7	18.5	23.9	24.9	26.0
Current liabilities	17.0	17.7	18.5	23.9	24.9	26.0
Long-term debt	2,668.4	2,771.4	2,871.7	3,037.4	3,147.0	3,251.8
Total liabilities	2,685.4	2,789.1	2,890.1	3,061.3	3,171.9	3,277.8
Shareholders equity	607.7	379.9	170.9	490.2	269.2	70.3
Total liabilities & shareholder equity	3,293.2	3,169.0	3,061.1	3,551.5	3,441.1	3,348.1

Figure 12.10: Market and Book Value Balance Sheets (\$000s)

P&L (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
Sales	2,000.0	2,137.5	2,285.6	2,445.3	2,617.5	2,803.2
Expenses	584.2	605.2	627.1	760.4	789.6	820.1
EBITDA	1,415.8	1,532.3	1,658.5	1,684.9	1,827.8	1,983.1
Depreciation	185.0	174.3	161.4	149.4	175.4	162.4
EBIT	1,230.8	1,358.0	1,497.2	1,535.5	1,652.5	1,820.7
Interest	99.0	146.8	152.4	157.9	167.1	173.1
EBT	1,131.8	1,211.3	1,344.7	1,377.5	1,485.4	1,647.6
Taxes	0.0	0.0	0.0	0.0	0.0	0.0
Earnings	1,131.8	1,211.3	1,344.7	1,377.5	1,485.4	1,647.6

Figure 12.11: Pro Forma P&Ls (\$000s)

by backing it out from the balance sheet identity, equation 12.10. However, we can also estimate the equity value by discounting the equity free cash flows by the equity cost of capital. The equity cost of capital follows from equation 11.5, or 11.4 because there is no excess cash, in chapter 11.

$$r^E = r^O + \frac{Debt - Cash}{Equity} (r^O - r^D) = 0.0941 + 0.11 \times (0.0941 - 0.0550) = 9.85\%$$

The levered cost of capital is a bit higher than the unlevered cost of capital precisely because Bill maintains debt in his capital structure. However, the difference between the

Free cash flow (\$000)	2021A	2022E	2023E	2024E	2025E	2026E
NOPAT (EBIT x (1-t))		1,358.0	1,497.2	1,535.5	1,652.5	1,820.7
Depreciation		174.3	161.4	149.4	175.4	162.4
Capex		0.0	0.0	500.0	0.0	0.0
Change in NWC		49.4	52.7	139.1	64.2	68.5
Unlevered free cash flow		1,482.9	1,605.9	1,045.8	1,763.7	1,914.6
Present values		1,355.3	1,341.4	798.4	1,230.6	1,221.0
Sum (short-term value)	5,946.8					
After-tax interest expense		146.8	152.4	157.9	167.1	173.1
Net debt issuance		102.9	100.3	165.7	109.6	104.8
Levered free cash flows		1,439.0	1,553.7	1,053.6	1,706.2	1,846.3
Present values		1,310.0	1,287.6	794.9	1,171.8	1,154.3
Sum (short-term value)	5,718.7					

Figure 12.12: Unlevered and Levered Free Cash Flows (\$000s)

two expected returns is relatively small, $9.85 - 9.41 = 0.44\%$, because on a market value basis Bill only funds 10% of the his business' value with debt. Also, notice that we don't need to know the equity value - which is what we're trying to find - to estimate the equity cost of capital because we have assumed Bill maintains a target (net) leverage ratio of 10%. Therefore, the ratio of debt to equity is $0.1/0.9 = 0.11$.

Discounting the equity free cash flows by the equity cost of capital produces the short-term value of equity.

$$\begin{aligned} \text{Short-term value of equity in year 0} &= \frac{1,439,036}{1 + 0.0985} + \frac{1,553,743}{(1 + 0.0985)^2} + \dots + \frac{1,846,332}{(1 + 0.0985)^5} \\ &= \$5,718,665 \end{aligned}$$

Estimation of the long-term or terminal value of equity follows immediately from our estimate of the value of the firm and assumed target leverage ratio. The terminal value of the firm is \$32,517,635. Bill plans on maintaining a 10% target leverage ratio implying a terminal equity value of $(1-0.10) \times 32,517,635 = \$29,265,872$. Discounting this amount back to year zero produces $29,265,872 / (1 + 0.0985)^5 = \$18,297,292$. Adding this to the short-term value of the firm gives us a total equity value today equal to $5,718,665 + 18,297,292 = \$24,015,957$. This is exactly what obtained from backing out the equity value using the estimated enterprise and debt values.

12.2.6 Financial Statement Analysis

Figure 12.13 presents a financial statement analysis of Bill's pro forma financial statements. (See the Accounting Appendix for definitions of these different financial ratios.) Performing

financial statement analysis on our valuation model serves several purposes. First, it provides a reality check on our model. Are the implied margins, returns, and other metrics plausible? Second, it helps in identifying future risks. Will credit risk become unsustainable? Will we violate a loan covenant? Third, it enables us to benchmark our expected performance against our peers and competitors who make available financial information (e.g., publicly traded companies).

Financial statement analysis	2021A	2022E	2023E	2024E	2025E	2026E
P&L metrics						
Sales growth		6.9%	6.9%	7.0%	7.0%	7.1%
Operating income growth		10.3%	10.2%	2.6%	7.6%	10.2%
Margins						
EBITDA	70.8%	71.7%	72.6%	68.9%	69.8%	70.7%
EBIT	61.5%	63.5%	65.5%	62.8%	63.1%	65.0%
Net	56.6%	56.7%	58.8%	56.3%	56.8%	58.8%
Returns						
ROA		42.0%	48.1%	46.4%	47.3%	53.6%
ROIC		42.3%	48.3%	46.7%	47.6%	54.0%
ROE		245.3%	488.2%	416.7%	391.2%	970.6%
DuPont						
Net margin (Earnings/Sales)		0.57	0.59	0.56	0.57	0.59
Asset turnover (Sales/Assets)		0.66	0.73	0.74	0.75	0.83
Leverage (Assets/Equity)		6.54	11.31	10.00	9.21	20.00
ROE		245.3%	488.2%	416.7%	391.2%	970.6%
Working capital						
Days cash	270	270	270	270	270	270
Days receivable	90	90	90	90	90	90
Days payable	15	15	15	15	15	15
Cash conversion cycle	75	75	75	75	75	75
Credit metrics						
EBITDA-to-interest	14.3x	10.4x	10.9x	10.7x	10.9x	11.5x
Debt-to-EBITDA	1.9x	1.8x	1.7x	1.8x	1.7x	1.6x
Current ratio	55.2x	55.7x	56.4x	61.3x	62.1x	62.8x
Acid ratio	55.2x	55.7x	56.4x	61.3x	62.1x	62.8x
Market leverage ratio	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%

Figure 12.13: Financial Statement Analysis

Top line growth increases slightly over time due to the increasing share of revenue coming from cosmetic procedures. Operating income (EBIT) growth shows a significant drop in 2024 when Bill brings on a new hygienist, but it quickly rebounds as a result of his strong earnings margins. Bill runs a highly profitable business as suggested both by his large margins and high returns. His ROIC increases from 42% to over 50% by 2026, suggesting that for each

dollar of invested capital - debt and book equity - Bill will generate nearly 54% by 2026.

Bill's return on equity is astronomical, reaching almost 1000% by 2026. The DuPont decomposition reveals why it's so large. While he has both a healthy margin and asset turnover, Bill's financial leverage is quite high. He relies mostly on debt financing for the business, using most of his equity free cash flows for his personal income.

This large amount of debt relative to the total invested capital may be cause for concern. However, a look at the credit metrics reveal that, in fact, Bill's business is fairly conservatively financed. His interest coverage ratio is always above 10 implying for each dollar of interest expense Bill has over \$10 of cash earnings. Similarly, his debt-to-EBITDA ratio is always under two suggesting that he only needs to year's worth of EBITDA to pay off all of his debt. His liquidity position is also reassuring. For each dollar of current liabilities, Bill has over \$50 of current assets - cash and receivables.

The dark side to this amount of liquidity is that Bill is not managing his working capital very efficiently. It takes almost three months for his customers pay him while he pays in employees every two weeks. As a result, he must maintain a relatively large cash balance to ensure against cash shortfalls. But, holding a large cash balance is inefficient. If Bill could get his customers to pay within 60 days, for example, that would free up more cash for investment purposes and increase the value of his business. It would also likely alleviate some risk from having to wait so long to collect money.

12.3 Market Multiples and Comparable Companies

Whereas DCF valuation provides an **intrinsic value** based on assumptions about the companies fundamental business drivers, **market multiples** valuation provides a **relative value** using the values of peer or **comparable companies**, so-called "**comps**".. Before performing a multiples valuation on Bill's dental practice, let's provide some intuition with a real estate example.

12.3.1 Valuing a Home

Imagine we want to sell our house, and we're not sure of the price. Our home is located in Ardmore, Pennsylvania, is 3,000 square feet in size, and has 4 bedrooms and 3.5 bathrooms. A common approach to estimating the price employs the following recipe.

1. Find comparable homes. Comparable means similar location, school district, number of bedrooms/bathrooms, amenities (large backyard, pool, close to public transportation, etc.), etc. Imagine we found five comparable homes, which we'll refer to as homes A, B, C, D, and E.
2. For each comparable home compute a price per square foot, which is an example of a **market multiple** or more simply **multiple**. The price could come from a recent transaction, or from a home appraiser or listing service like Zillow or Redfin.⁴ The table below presents the price, square footage, and market multiple for each of our comparable homes.

House	Price	Square feet	Price/ft ²
A	825,000	3,200	257.8
B	1,275,000	3,900	326.9
C	984,000	3,000	328.0
D	1,105,000	4,100	269.5
E	1,098,000	3,650	300.8

3. Aggregate the multiples for the comparable companies by taking an average or median. The average and median price per square foot of our five homes are \$296.6 and \$300.8, respectively.
4. Multiply the aggregated comps multiple by our home's square footage to get our home's price estimate.

$$Price_{\text{Our home}} = \underbrace{\left(\frac{\$296.6}{ft^2} \right)}_{\text{Comps multiple}} \times \underbrace{3,000 ft^2}_{\text{Our home } ft^2} = \$889,842$$

While mathematically market multiples valuation is a trivial exercise, the devil is in the details. The accuracy of our estimate depends in large part upon the quality of the comparable homes. The more similar the comparable homes are to our home, the better the estimate. However, this does not mean that the only valid comparable homes are those that are identical to our home. Such a home doesn't exist. Even in a development with identically constructed homes, differences in interior finishings and location (near the entrance or common area) can affect the price of the home.

⁴When valuing businesses, we distinguish comparable companies whose price is based on a recent transaction vs. current market price.

How Many Comparables?

The number of comparable homes we use to estimate the value of our home reflects a tradeoff. The more homes we consider, the more statistically precise our home estimate is but at the expense of including less similar homes and introducing bias into our estimate. In statistics, this is the classic bias-variance tradeoff. Let's explain by consider two extremes.

We could just look at the price of the one home most similar to ours? For example, home C is the same square footage and let's assume it has the same number of bedrooms and bathrooms and is just down the street from our home and has the same attached two car garage etc. If home C is more similar to our home in every observable dimension, why bother with other homes? We could just use the price per square foot of home C to estimate the price of our home?

Despite these similarities, we can still wind up with a very poor estimate of the value of our home, far worse than one that incorporates information from other, less similar homes. To see how, consider the following situations. Home C's owners were in a financially precarious position when they sold the home and took the first low-ball offer they received. Or, the buyers of home C were relatives of the owners and received a substantial "family discount." Or, the buyers of home C were relatives of the neighbors and offered an exorbitant amount to be close to family. Each of these situations will lead to home prices that have less to do with the home and more to do with features of the buyer or seller. Ultimately, we can get a very poor estimate of the value of our home despite the close similarity to the Home C.

Alternatively, we could look at the price of every home in the same school district? Averaging over such a large number of data points will surely give a very precise estimate of our home's value. The problem is that it may also give a very biased estimate of our home's value. Imagine that our home is situated in a very desirable location within the school district - walking distance to the school, close to public transportation, etc. Buyers pay a premium for this unique location that will be underweighted in the grand average of all homes in the district. In other words, we'll wind up with an estimated that is biased down and underestimates the value of our home.

So, the goal in forming a list of comparable homes is not to find the single best match nor to use as many homes as possible. The goal is to find comparable that are similar to ours along the most important dimensions - neighborhood, size, bedrooms/bathrooms, etc. - and that differ from our home in different ways. For example, home B has nicer finishings - marble countertops, new hardwood floors, etc. - than our home. Home A has one fewer bedroom. Home D is an extra half bathroom, while home E doesn't have a garage. Each

home is somewhat different from our home but in different ways. So, together the portfolio of homes is plausibly more similar to our home than any individual home.

An analogy for comparable homes, and companies, comes from Michael Lewis' book, *Moneyball* (great book), which was made into a movie (good movie). I'll take a little liberty with the details to make the point, but the story goes like so. The Oakland A's baseball team had just lost a star player, Jason Giambi, who moved to the New York Yankees. In trying to replace Giambi, the A's manager, Billy Beane, realized they couldn't find a single player like him, but they may be able to find several players whose combined attributes - on-base percentage, fielding, run-batted-in - replicated Giambi's. We want to find comparable homes that together contain the features of our home but differ from our home in different ways.

Ultimately, the exact optimal number of comps is not obvious, and would likely vary from situation to situation. However, between 10 and 20 comps is not an uncommon range.

Market Conditions

By using market prices from other homes as the basis for valuing our home, the estimated value embeds current market conditions. If the housing market is particularly hot, valuations will be elevated and the implied price of our home will be elevated. If the housing market is particularly cool, valuations will be depressed and the implied price of our home will be depressed. Depending on whether we're buying or selling, this feature of multiples valuation can be comforting or troubling. Regardless, this is a key feature of all multiples valuations.

12.3.2 Conceptual Framework

Valuing a company with market multiple is conceptually no different from valuing a home. We follow the same recipe. We start with a target firm we want to value. Let's call it firm "T," for target. First, we need to come up with a list of companies that are comparable to firm T *and* have an observable market value. This means comparable companies must be publicly traded. Second, we need to compute a market multiple for each of our comparable companies. Third, we need to aggregate the comparable company multiples by taking an average or median or some other statistic. Finally, we can apply this aggregate multiple from our comparable firms to our target firm to get an estimate of value.

But, this recipe in the context of companies raises several questions. What market multiple should we use? What does it mean for a firm to be comparable? This section

answers these questions by starting from first principles - the Fundamental Value Relation - and prepares us to value VB dental using market multiples.

It's easy to think that multiples valuation is unrelated to discounted cash flow analysis and our Fundamental Value Relation. It's not. Consider our target firm which has free cash flows, CF^T , that grow at a constant rate g and have a cost of capital, r . The value of firm T can be written expressed as a growing perpetuity.

$$Value^T = \frac{CF^T}{1+r} + \frac{CF^T(1+g)}{(1+r)^2} + \frac{CF^T(1+g)^2}{(1+r)^3} + \dots = \frac{CF^T}{r-g}$$

Now consider another firm "C," for comparable, with a different cash flow, CF^C , but the same growth rate g and cost of capital r as our target firm. Firm C's value can also be expressed as a growing perpetuity.

$$Value^C = \frac{CF^C}{1+r} + \frac{CF^C(1+g)}{(1+r)^2} + \frac{CF^C(1+g)^2}{(1+r)^3} + \dots = \frac{CF^C}{r-g}$$

Because the firms have different cash flows, $CF^T \neq CF^C$, they have different values, $Value^T \neq Value^C$. But, consider their value multiples by dividing both sides of the two expressions by their respective cash flows.

$$\frac{Value^T}{CF^T} = \frac{Value^C}{CF^C} = \frac{1}{r-g} \quad (12.11)$$

The value multiples of the two firms are equal and a function of (1) the discount rate r and (2) the cash flow growth rate g . This result has three key implications.

First, valuation by market multiples is not fundamentally different from discounted cash flow analysis. Market multiples valuation is derived from our Fundamental Value Relation. What differs from DCF are the inputs. Second, equation 12.11 provides a foundation for market multiples valuation. Imagine we don't know the value of firm T, $Value^T$. We can find a comparable firm, firm C, and use it's value multiple in concert with firm T's cash flow to estimate the value of firm T like so.

$$\frac{Value^T}{CF^T} = \frac{Value^C}{CF^C} \implies Value^T = \underbrace{\frac{Value^C}{CF^C}}_{\text{Comps multiple}} \times CF^T \quad (12.12)$$

Finally, equation 12.11 identifies the characteristics that comps should share with the target firm. When using cash flow multiples to estimate value, good comps have cash flows whose growth rate and systematic risk are similar to that of the company we are trying to value, i.e., the target.

Market Multiples and Implications for Comparable Firms

Table 2 displays some multiples commonly used in practice. The left column lists the different multiples, which are grouped according to the measure of value, i.e., the numerator of the multiple. Enterprise or firm value multiples are scaled by measures of the money available to or invested by *all* investors - debt and equity. Equity multiples are scaled by measures of the money available to or invested by only equity investors. This consistency between the numerator and denominator in valuation multiples is important in ensuring that the latter accurately captures as many key value drivers of the former.

Enterprise/firm value multiples	Comps characteristics to match
Unlevered free cash flow	Risk (r), growth (g)
NOPAT	+ investment requirements
EBIT	+ tax cost structure
EBITDA	+ depreciation cost structure
Sales	+ operating cost structure
Invested capital	(Similar to sales)
Equity value multiples	
Levered free cash flow	Risk (r^E), growth (g)
Earnings	+ investment requirements
Book equity	(Similar to sales)

Table 2: Market Value Balance Sheet - VB Dental 2021. Values in \$000s

The right column identifies the firm characteristics on which we want to match target and comparable firms. For example, when we use a cash flow multiple - unlevered or levered - we want to ensure, as much as possible, that the risk and growth of the comparable firms cash flows matches those of our target firm's cash flows. This logic follows from equation 12.11.

When using a NOPAT multiple, the cash flow risk, growth, *and* investment policies of our comparable firms should match those of our target firm. When we use an EBIT multiple, the cash flow risk, growth, investment policies, *and* tax cost structure of our comparable firms should match those of our target firm. And so on.

In other words, the characteristics on which we match depends on the multiple we're using to value our target firm. The further removed our multiple's denominator is from free cash flow, the more features on which we have to match. To understand this logic, consider NOPAT and our free cash flow definition in equation 12.7 above. The difference between

NOPAT and free cash flow is that the latter accounts for investment in long-term assets and working capital. Any differences in investment policy between our target and comparable firms won't affect NOPAT but will affect value leading to different multiples. Thus, we have to match on risk, growth, and investment requirements.

Selecting Comparable Firms

A natural starting point for finding firms comparable to our target firm is the industry of the target firm. Industry is a catch-all proxy for a variety of value relevant firm characteristics, such as risk and growth. Industry is typically defined by Standard Industrial Classification (SIC) code, North American Industry Classification System (NAICS) code, or Global Industrial Classification code (GICS). All three measures are similar in spirit so we'll focus on SIC codes.

SIC codes are 4-digit, numeric codes whose interpretation according to the Occupational Safety and Health Administration is provided in the table below. Each additional digit in the SIC code narrows the classification of firms. Table 3 describes this hierarchy and provides an example.

SIC Code		
Digits	Category	Example
1st	Divisions	5 - Wholesale & Retail Trade
1st - 2nd	Major group	57 - Home Furniture, Furnishings, And Equipment Stores
1st - 3rd	Industry group	571 - Home Furniture And Furnishings Stores
1st - 4th	Specialty group	5713 - Floor Covering Stores

Table 3: SIC Code Hierarchy

Selecting firms from the same SIC group as the target comes with challenges as suggested by Table 4. As we move from 1-digit to 4-digit classifications, two things happen. First, the number of firms in a group rapidly declines. The median number of firms in a 2-digit major group is 23 compared to 7 for a 3-digit industry. Second, the comparability among firms in groups increases. So, while 3-digit industries have fewer firms from which to select comparables, they tend to be more alike. Using the example from Table 3 above, the specificity of the business is increasing as we move from 1- to 4-digit SIC code. So, there's a tradeoff. A narrower industry definition may ensure greater comparability but at the expense of fewer comparable firms.

Additionally, selecting comparable firms from the same industry is often insufficient. We must often match on other characteristics including: firm size (market capitalization,

	Division (1-Digit)	Major Group (2-Digit)	Industry Group (3-Digit)	Specialty Group (4-Digit)
Count	10	71	258	406
Average	802	113	31	20
Minimum	40	1	1	1
25th percentile	207	12	3	2
Median	354	23	7	5
75th percentile	918	64	15	11
Maximum	3,912	3,144	2,812	2,635

Table 4: SIC Code Hierarchy

sales total assets), projected growth (revenue, assets), geographic location of sales, product market, business model (subscription vs. transaction), regulatory and legal environment, growth strategy (organic vs. acquisitions), etc. Of course, the more characteristics on which we match, the fewer comparable firms with which we are left. So, the goal isn't to exactly match firms along all dimensions. Instead, we want to be aware of the dimensions on which our comparable firms are close or distant matches to our target firm to ensure that there are no gaps.

Sometimes comparable firms come from different industries. Consider Ferrari Inc. (RACE), which went public on October 21, 2015. Ferrari's revenue comes primarily from manufacturing and selling luxury vehicles. It also sells car engines to other auto manufacturers (e.g., Maserati) and manages a Formula 1 race team. As an auto manufacturer, we might instinctively think of comparable firms as other auto manufacturers. But, I don't think anyone is going to confuse a Ferrari with a Ford. Put differently, while both Ferrari and Ford manufacture and sell cars, their business models, target markets, risks, size, growth, operations, etc. are very different. So we might narrow the set of auto manufacturers to brands like Lamborghini, Porsche, Aston Martin, Bentley, and Rolls Royce. However, as of 2015, none of these companies were publicly traded. Most are brands owned by larger auto manufacturers such as Volkswagen and BMW. Instead, we may want to compare Ferrari to other luxury brands, such as Hermes, LVMH, and Prada. While the products these companies make - clothing, personal accessories, liquor - are quite different from cars, their target markets, risks, and other aspects are very similar. This is all to say that comparable companies are defined by similar value drivers, not just similar products.

12.3.3 Valuing VB Dental

Now let's value VB Dental using market multiples. Finding a publicly traded dental office run by one dentist out of Florida is going to be impossible. In fact, finding a publicly traded dental office in the U.S. run by any number of dentists will be all impossible, at least as of 2024. However, remember that the goal with comparable companies is not to find an exact match of the target. Rather, it's to find a sufficiently larger number of companies that are similar with regards to fundamental value drivers. Precisely which drivers depends on the multiple we'll be using; however, in practice we should consider the effects of all drivers - risk, growth, taxes, investment, cost structure, etc.

To manage the discussion and analysis, we'll focus on five comparable companies for VB Dental, which are listed in Table 5. Five is on the (too) low side in term of number of comparables, but our goal here is to clearly communicate the process. From the descriptions, none of the companies are dental, or even medical, practices. They are nonetheless members of the dental supply chain. We could have also considered physician practice management companies that focus on running the business, as opposed to care, aspects of a medical practice. Health insurance companies would probably not be a good choice because of their distinct business model and regulatory requirements.

Company	Description
Align Technology, Inc. (ALGN)	Global medical devices for teeth straightening.
Henry Schein, Inc. (HSIC)	Dental and medical supply company
Patterson Companies, Inc. (PDCO)	Dental and veterinary products and services
Envista Holdings Corporation (NVST)	Global family of more than 30 dental brands
Dentsply Sirona Inc. (XRAY)	Dental equipment manufacturer

Table 5: SIC Code Hierarchy

For each comparable company, we require financial data to estimate their market multiples. This data are presented in Table 6. A couple of comments about this comparable company data. First, the data should be contemporaneous, or as close as possible in time, with that of the target company. Current equity values are relatively easy to find or estimate given the liquidity of equity markets. Accounting data is reported less frequently. This means using the **trailing twelve months (TTM)** of data, which can be assembled from quarterly (10-Q) and annual (10-K) reports. For example, if we are valuing a company as of July 31 we would ideally like information for comparables as of July. As most companies have December fiscal year ends, this mean assembling data as of June 30. For

companies with other fiscal year ends, this may mean assembling data as of April 30 or May 31. This gap between our valuation date and the date of the comparable companies data means we should peruse other information sources falling in this gap, such as 8Ks and company announcements, for changes to this data.

	Comparable Company					Average	Median
	ALGN	HSIC	PDCO	NVST	XRAY		
	Value						
Cash	898.9	159.0	114.5	948.5	291.0		
Debt	95.1	2,718.0	763.0	1,646.0	2,210.0		
Equity	18,298.0	8,231.0	2,109.0	2,831.0	5,240.0		
Enterprise value	18,393.1	10,949.0	2,872.0	4,477.0	7,450.0		
	Value Driver						
Sales	3,916.5	12,451.0	6,568.3	2,562.9	3,940.0		
EBITDA	816.8	919.0	341.1	424.9	619.0		
EBIT	677.3	650.0	252.9	265.8	274.0		
NOPAT	471.3	505.1	193.0	265.8	274.0		
Earnings	462.3	388.0	185.9	(120.4)	(95.0)		
	Multiple						
EV-to-Sales	4.7	0.9	0.4	1.7	1.9	1.9	1.7
EV-to-EBITDA	22.5	11.9	8.4	10.5	12.0	13.1	11.9
EV-to-EBIT	27.2	16.8	11.4	16.8	27.2	19.9	16.8
EV-to-NOPAT	39.0	21.7	14.9	16.8	27.2	23.9	21.7
Price-to-Earnings	39.6	21.2	11.3	N/A	N/A	24.0	21.2

Table 6: VB Dental Comparable Company Financial Data (\$million) and Multiples

We also need to scrub or clean our comparable company data of any anomalies or transitory shocks. For example, imagine a temporary increase in costs or the exhaustion of historical tax shields. Both of these shocks may have a significant impact on current earnings measures. However, because they are temporary, they are unlikely to have a significant effect on value, which is a function of *all* future earnings. As a result, the comparable company's multiple will be affected, which will affect the valuation of the target. In these circumstances, we would remove the transitory shock from the value driver by using a projected sales or earnings or historical average.

Having scrubbed the comparable company financial data and computed their multiples, we can now value our target firm by multiplying the comparable company multiples by the target firm's value driver (sales, EBITDA, EBIT, NOPAT, and earnings). Table 7 presents the implied market multiple valuations of VB Dental. The top half of the table

presents estimated enterprise values, the bottom half presents estimated equity values. The first columns presents VB Dentals' corresponding value driver, i.e., Sales, EBITDA, EBIT, NOPAT, and Earnings.

Multiple	VB Dental Value Driver	Comparable Company					Average	Median
		ALGN	HSIC	PDCO	NVST	XRAY		
Enterprise Values								
EV-to-Sales	2.0	9.4	1.8	0.9	3.5	3.8	3.9	3.5
EV-to-EBITDA	1.4	31.9	16.9	11.9	14.9	17.0	18.5	16.9
EV-to-EBIT	1.2	33.4	20.7	14.0	20.7	33.5	24.5	20.7
EV-to-NOPAT	1.2	48.0	26.7	18.3	20.7	33.5	29.4	26.7
Equity Values								
EV-to-Sales		6.7	N/A	N/A	0.8	1.1	1.2	0.8
EV-to-EBITDA		29.2	14.2	9.3	12.2	14.4	15.9	14.2
EV-to-EBIT		30.8	18.1	11.3	18.1	30.8	21.8	18.1
EV-to-NOPAT		45.4	24.0	15.6	18.1	30.8	26.8	24.0
Price-to-Earnings	1.1	44.8	24.0	12.8	N/A	N/A	27.2	24.0

Table 7: VB Dental Market Multiples Valuations (\$millions)

For example, the enterprise value implied by Align Technology's enterprise value-to-EBITDA multiple is $22.518 \times 1.416 = \$31.9$ million. Similarly, the equity value implied by Patterson Companies price-to-earnings ratio is $11.343 \times 1,132 = \$12.8$ million. We obtain equity values from from estimated enterprise values by subtracting debt (\$2.67 million) and adding excess cash (\$0).

Clear from Table 7, the estimated valuations vary a great deal. Enterprise values range from \$0.9 million to \$48 million. Equity values range from \$0.8 million to \$45.4 million. Focusing on averages or medians mitigates some of the variability. However, both average and median enterprise and equity values still vary a great deal across different multiples. This is common in practice.

Which one is correct or most accurate? It's unclear, though we know from our discussion above that multiples with value drivers closer to free cash flow require fewer dimensions on which to match and therefore may leave less room for bad matches. This suggests that EV-to-NOPAT or perhaps price to earnings may be the most reliable and that sales multiples should be treated with caution. In practice, EV-to-EBITDA multiples are arguably the most commonly used though the theoretical justification for this emphasis is also unclear.

Perhaps more interesting is how the multiples valuations relate to the DCF valuation. DCF analysis produced an enterprise value estimate of \$26.7 million, almost identical to the

valuation implied by the median enterprise value-to-NOPAT multiple. More generally, the DCF estimate is higher than most market multiple implied valuations. There are several potential explanations for this relation.

1. Dental/medical stock valuations may be depressed because the market has taken a pessimistic view of their future potential.
2. Our comparable companies are poor matches to our target firm in that their cash flow growth may be lower or risk may be higher.
3. Our DCF cash flow forecasts are overly optimistic, or our estimated cost of capital is too low.

Any or all three explanations may be behind any discrepancy between the DCF and market multiples estimates. Nonetheless, the combination of DCF and market multiple estimates are both informative because they rely on different assumptions about value drivers. DCF relies on our, the analysts, assumptions. Multiples rely on the market's assumptions.

12.4 Key Ideas

- 1.

12.5 Technical Appendix

Appendix A

Personal Budgeting

This appendix discusses personal budgeting. While unnecessary to understand any chapter in the book, budgeting provides additional context and, more importantly, is a critical life skill. Without a budget, there is no way to know what we're spending and where we're spending it, or whether we're saving enough. The consequences of not following a budget can be dire. In the short term, we could run out of money for rent, groceries, or other basic needs. In the long-term, we could run out of money in retirement.

Not only can a budget help us avoid these mistakes, but it can inform our decision making. A budget details how much money we get, where we spend it, and what, if anything, is left over to save. The best part of a budget is that once set up, managing it requires little effort but provides enormous benefit.

Like all chapters, this appendix is accompanied by an Excel spreadsheet containing our example budget, with a few bells and whistles. It can be used as a starting point for your own budget.

A.1 Setting up a budget

Conceptually, a budget is simple. For each period, often a pay-period (bi-weekly, monthly) over a horizon (one year), we need to estimate two things.

1. **Cash inflows** or **Income**. The money we receive each pay period.
2. **Cash outflows** or **Expenses**. The money we spend each pay period.

The difference between 1. and 2. are our savings when our income is greater than our expenses, or what we need to borrow if our income is less than our expenses. (By the way, this budgeting process is also exactly what companies do. Sorry...I can't resist emphasizing that there is only "one" set of financial concepts for everyone and everything.)

Let's illustrate the process by creating a budget for Leslie, a 34-year old financial executive with two kids and a stay at home husband. Leslie and her family rent an apartment in New York.

A.2 Income - Cash Inflows

A.3 Salary

For most, estimating income is simple because it's the same amount each pay period for the year. In other words, we can just look at our paycheck or direct deposit in our bank account and see how much money we received. For others, estimating income can be a little tricky because they don't receive the same amount every pay period. Many executives receive bonuses contingent on the performance of their company, and these bonuses can come in multiple forms, such as cash, stock, and options. Likewise, many sales people often work on commission, meaning they get a fraction of what they sell. If they sell a lot one month, they make a lot. If they don't, they make less. Nonetheless, experience and expectations can help these employees estimate how much money they might receive and when.

Leslie's salary is \$480,000 paid monthly, so \$40,000 per month. She is also eligible for a bonus in December equal to one to three times her monthly salary, so \$40,000 to \$120,000. Of course, this is *not* what Leslie actually receives because these numbers correspond to her **gross salary**, where gross in this context means before any deductions (e.g., taxes, insurance) that the company makes on behalf of Leslie.

1.3.1 What is (Not) in My Paycheck?

Paychecks come in a variety of different forms and the information on them varies widely. That said, there is some information that is generally common to all paychecks. Figure A.1 presents Leslie's March salary statement. Leslie uses direct deposit so rather than receiving a check, her employer deposits her paycheck directly into her bank account.

Large Financial Institution 1 Wall Street New York New York, 11001.
Leslie Ann Palmer Nice Street in West Village, Unit PH4 New York, New York 11004

Name	Company	Employee ID	Pay Period Begin	Pay Period End	Check Date	Check Number
Leslie Ann Palmer	Large Financial Institution	024731-934	03/01/2023	03/31/2023	03/31/2023	

	Gross Pay	Pre Tax Deductions	Employee Taxes	Post Tax Deductions	Net Pay
Current	40,000	2,669.58	16,469.65	304.06	20,556.71
YTD	120,000	8,008.74	49,489.85	912.18	61,589.23

Earnings						Employee Taxes		
Description	Dates	Hours	Rate	Amount	YTD	Description	Amount	YTD
Regular Monthly Pay	03/01/2023 - 03/31/2023	0	0	40,000.00	120,000.00	OASDI	2,314.48	6,943.48
						Medicare	541.29	1,623.87
						Federal Withholding	9,580.79	28,742.37
						State Tax - NY	2,557.13	7,671.40
						SUI-Employee Paid - NY	29.02	167.97
						City Tax - Manhattan	1,446.93	4,340.78
Earnings				40,000.00	120,000.00	Employee Taxes	16,469.65	49,489.85

Pre Tax Deductions			Post Tax Deductions		
Description	Amount	YTD	Description	Amount	YTD
Dental Big Insurance Co.	82.03	246.09	Supplemental LTD	251.91	755.73
Medical Great Insurance Policy	693.00	2,079.00	Supplemental Life Insurance	52.15	156.45
Employee Retirement	1,875.00	5,625.00			
Vision Plan	19.55	58.65			
Pre Tax Deductions	2,669.58	8,008.74	Post Tax Deductions	304.06	912.18

Employer Paid Benefits			Taxable Wages		
Description	Amount	YTD	Description	Amount	YTD
Basic Life Insurance Employer	28.28	84.84	OASDI - Taxable Wages	40,000.00	120,000.00
Dental Big Insurance Co.	46.40	139.20	Medicare - Taxable Wages	40,000.00	120,000.00
Employer Retirement Basic	1,875.00	5,626.00	Federal Withholding - Taxable Wages	40,000.00	120,000.00
Employer Retirement Match	1,875.00	5,626.00	State Tax Taxable Wages - NY	40,000.00	120,000.00
Medical Great Insurance Policy	1,246.01	3,738.03	City Tax Taxable Wages - Manhattan	40,000.00	120,000.00
Employer Paid Benefits	5,070.69	15,212.07			

	Federal	State	Absence Plans			
Marital Status	Married		Description	Accrued	Reduced	Available
Allowances	2	0	Advanced Sick Time	0	0	178
Additional Withholding	0	0				

Payment Information					
Bank	Account Name	Account Number	USD Amount	Amount	
Employee Direct Deposit	Employee Direct Deposit *****0469	*****0469		20,556.71	USD

Figure A.1: Leslie’s Paycheck Information

There’s a lot of information in the statement. Fortunately, it is organized into groups. The top box provides information about the employee (Leslie and her Employee ID), employer (Large Financial Institution) and the pay period (03/01/2023 to 03/31/2023). The next box summarizes Leslie’s salary and deductions for the current pay period (March 2023) and year-to-date or YTD (01/01/2023 to 03/31/2023). In March 2023, Leslie’s gross pay was \$40,000. From that gross pay, her company made three types of **deductions** (i.e., subtractions): pre-tax, tax, and post-tax. As the names suggest, pre-tax deductions are removed from her gross salary before taxes are removed. Post-tax deductions are removed from her salary after taxes are removed. The relevance of this distinction will become clear a moment.

Earnings

Leslie’s earnings are repeated in the Earnings box. Some employees may have more than one source of income from their employer, for example, a base salary plus a commission. These different sources are usually shown here and would add up to the gross pay shown at the top.

Pre-Tax Deductions

Leslie's pre-tax deductions are listed next on the left side. Leslie has chosen to have these deductions from her paycheck, i.e., they are optional. These deductions include: dental, health, and vision insurance for her and her family, as well as a contribution to her employer's retirement plan. This last deduction is not a gift to the company. Rather, it's a contribution to a savings account that her employer has set up for her.

In March, \$2,669.58 goes towards these pre-tax deductions. This means that the income used to calculate her taxes is $40,000 - 2,669.58 = \$37,330.42$. This lower income will reduce the taxes she pays as we'll see in a moment.

Employee Taxes

Employee taxes are detailed next. Leslie pays six different types of taxes.

1. **OASDI.** Taxes you pay towards social security, the government pension for retirees.
2. **Medicare.** Taxes you pay towards Medicare, the government sponsored health insurance for person 65 and older and some younger people with disabilities.
3. **Federal Withholding.** Taxes you pay to fund the operation of the federal government.
4. **State Tax - NY.** Taxes you pay to the state, in this case New York, to support the operation of the state government.
5. **SUI-Employee Paid - NY.** Taxes you pay towards the state unemployment insurance that aids unemployed workers.
6. **City Tax - Manhattan.** Taxes you pay to the city, in this case Manhattan, to support the operations of the municipal government.

The application of these taxes varies greatly. Some states (e.g., Texas, Nevada, Washington) don't have a state income tax. Many cities (e.g., Anchorage, AL, Tampa, FL) don't impose an income tax. OASDI taxes only apply to the first \$160,200 as of 2023, meaning all income above that amount does not have to pay OASDI tax. The Medicare tax decreases from 1.45% to 0.9% after \$200,000 of income as of 2023. Thus, taxes, and by extension the money we actually receive, can vary throughout the year.

Because many taxes are computed as a fraction of income, pre-tax deductions can lower Leslie's taxes. For example, the OASDI tax rate is 6.2%. Without the pre-tax deductions, she would have owed $40,000 \times 0.062 = \$2,480$. Because of the pre-tax deductions, she owes $\$37,330.42 \times 0.062 = \$2,314.49$ - a savings of \$165.51. Similarly, the Medicare tax rate is 1.45%. The pre-tax deductions reduce her medicare taxes by $2,669.58 \times 0.0145 = \38.71 . Similar savings apply for her other taxes, especially federal withholding which is the largest.

Post-Tax Deductions

Leslie has only two post-tax deductions: supplemental long-term disability insurance and life insurance. These are additional insurance policies offered by her employer that supplement any other disability and life insurance policies that she may have.

Employer Benefits

Also detailed on Leslie's pay statement are all the benefits she receives that are paid by her employer. Her employer is contributing money towards her dental, health, and life insurance. They are also matching her retirement contributions. In other words, for each dollar that Leslie contributes to her 401-k savings plan, her employer contributes a dollar to her plan. That's additional money, no different from her salary except that it is not taxed until she withdraws it in retirement.

Taxable Wages

As mentioned above, different taxes are applied to wages differently. OASDI tax only applies to the first \$160,000 of earnings, meaning all money an employee earns above \$160,000 is not subject to OASDI tax. Some taxes, e.g., some city tax, apply to gross pay as opposed to pay less pre-tax deductions. This section tells Leslie the wages that were used to compute each tax.

Her estimated monthly gross salary for each month is presented in table 1. To fit on the page, we've collapsed January through November because the income and deductions are fairly similar.

To compute the money Leslie actually receives from her employer, we start with her gross salary and then subtract deductions. There are three types: pre-tax deductions, tax deductions, and post-tax deductions.

	Jan - Nov	Dec
Gross salary	40,000	120,000
Pre-tax deductions		
Health insurance	600	600
401-k contributions	1,875	0
Tax	14,822	41,193
Post-tax deductions		
Supplemental insurance	150	150
Total deductions	17,447	41,943
Net salary	22,553	78,057

Table 1: Leslie's Budget: Salary Component of Income (\$)

Pre-tax deductions get subtracted from the gross salary *before* the government computes how much we owe in taxes. Leslie's chosen two pre-tax deductions: health insurance and retirement savings. Each month, her company deducts \$600 from her paycheck to help cover the cost of the company's insurance plan that covers Leslie and her family. The company also deducts \$1,875 from her paycheck and puts it into an **employer-sponsored, defined contribution, personal pension account** or, more simply, a **401-k**.¹ Leslie's 401-k is basically a savings account in which she can choose, subject to some limitations, how to invest her savings (e.g., stocks, bonds).

Because her health insurance and 401-k contributions are deducted from her paycheck *before* the government computes the taxes Leslie owes, these deductions provide a **tax shield**. They reduce the taxes she has to pay to the government by lowering the amount of taxable income. When computing her taxes, rather than using \$40,000, the government uses $40,000 - (600 + 1,875) = \$37,525$. Because Leslie earns so much, she's in a very high tax bracket meaning she pays a large fraction of her income in taxes - 39.5% to be precise. By reducing her taxable income, the pre-tax deductions save Leslie $(600 + 1,875) \times 0.395 = \977.63 in taxes each month.

After pre-tax deductions come tax deductions, i.e., what the governments take. Governments is plural because federal, state, and sometimes local, governments take a piece of our salary. There are actually several different types of taxes (e.g., federal and state income tax,

¹Many non-profit organizations, such as universities, offer employees a 403-b, which is similar to a 401-k. Likewise, individuals can contribute money to an **individual retirement account** or **IRA**. There are, however, differences between these account types, such as how much we can contribute each year. This last point is important because there are strict limitations on how much individuals can contribute to these retirement plans each year. In Leslie's case, because she earns and contributes so much, she can only contribute to her 401-k for the first eight months of the year.

social security and medicare tax), but we've collapsed them into one line item in table 1. Finally, there may be post-tax deductions from a paycheck, such as the supplemental life insurance offered by her company to which Leslie has chosen to subscribe. Like the pre-tax deductions, these post-tax deductions are optional.

After all the deductions, Leslie's monthly gross income of \$40,000 has been reduced to \$22,553 of **net income** in January through November. (Ouch!!!!) This is the actual money she can spend or save - the cash inflow from her salary.

We might be wondering: Do we really need to lay out all those paycheck deductions - health insurance, 401-k contributions, etc.? No. We could have just recognized Leslie makes about \$22,500 each month, except for December when she receives her bonus.² For December we might estimate that her gross salary will be her base of \$40,000 plus an anticipated bonus of \$80,000 for total gross in come of \$120,000. At a 39.5% tax rate, this should leave her with \$72,600 of net income.

However, this quick and dirty approach hides how much Leslie is saving with her employer. That's her money and its an important component of her retirement savings plan. It also hides how much she's paying for benefits, which is important to know if she decides to switch jobs and wants to compare compensation packages. Finally, it hides just how much she's paying in taxes and why, especially at higher income levels, tax planning is so important.

A.4 Other Income

Some people have multiple income sources because they have multiple jobs or investment income. In Leslie's case, she only has one job, but she does receive income from two other sources of investment income.

1. **Securities income.** Leslie's existing savings generates income in the form of capital gains when she (or her mutual fund manager) sells assets for a profit, as well as dividends and interest income. Because she reinvests all of these earnings, she doesn't actually receive the money. Nonetheless, she has to pay taxes on these investment earnings.
2. **Rental income.** Leslie and her sister own an apartment that they rent. She receives half the rent, \$3,250, but budgets \$250 per month in maintenance and expenses

²Leslies after-tax income will vary January through November. In August, she will "max-out" her 401-k contributions in August and not be able to contribute anymore money. In April, she will have earned enough - \$160,000 as of 2023 - to fulfill her social security tax obligation and won't have to pay that tax anymore.

associated with the rental. So, her net, and taxable, income from the apartment is \$3,000.

Knowing all this additional income is important not just because it contributes to how much she can spend, in the case of rental income, but also because it is taxable. Keeping track of this income is imperative for not running afoul of the tax authority, or finding a large, unpleasant surprise in the form a big tax bill at the end of the year.

Table 2 estimates all of Leslie's income.

	Jan - Nov	Dec
Gross salary	40,000	120,000
Pre-tax deductions		
Health insurance	600	600
401-k contributions	1,875	0
Tax	14,822	41,193
Post-tax deductions		
Supplemental insurance	150	150
Total deductions	17,447	41,943
Net salary	22,553	78,057
Other income		
Securities income		20,000
Rental income	3,000	3,000
Total income	25,553	121,057

Table 2: Leslie's Budget: Total Income (\$)

Because the securities income is reinvested, Leslie has to recognize a corresponding expense - reinvested earnings - in December. She also has to recognize that come next April taxes on those investment earnings need to be paid. Depending on the tax status of the earnings, the tax bill could be as high as $20,000 \times 0.395 = \$7,900$.

1.4.1 Employee Retirement Savings Accounts

(Feel free to skip this section.) Leslie contributes to a traditional 401-k account in which *pre-tax money* - money deducted before the computation of taxes - is saved. She then chooses how this money is invested, subject to the investment options available in her companies retirement savings plan. Any growth in these savings arising from capital gains (selling assets

at a higher price), dividends (payments to stock owners), interest payments (payments to bond owners), are **tax deferred**, meaning Leslie won't have to pay any taxes on any investment income in her 401-k until she starts withdrawing money from the account. However, when Leslie retires and withdraws money from her 401-k, those withdrawals will be taxed as ordinary income, as if a company were paying her.

A ROTH 401-k (or ROTH IRA, ROTH 403-b) is an alternative employee retirement account. The key differences between the two are detailed in table 3.

	Traditional	ROTH
Contributions	Pre-tax meaning you pay less in taxes and have more money today	Post-tax meaning you pay higher taxes and have less money today
Investment earnings	Any investment earnings - capital gains, dividends, interest - are <i>not</i> taxed while the money stays in the account.	Same as Traditional
Withdrawals	No penalty for withdrawing money after age $59\frac{1}{2}$ but money is taxed as ordinary income at prevailing tax rate	No penalty for withdrawing money after 5 years and age $59\frac{1}{2}$ and money is tax-free

Table 3: Retirement Savings Accounts Features Summary

Table 4 shows Leslie's income breakdowns for a traditional and ROTH 401-k for comparison purposes.

Gross salary	40,000	Gross salary	40,000
Pre-tax deductions		Pre-tax deductions	
Health insurance	600	Health insurance	600
401-k contributions	1,875	401-k contributions	
Tax	14,822	Tax	15,563
Post-tax deductions		Post-tax deductions	
Supplemental insurance	150	Supplemental insurance	150
ROTH 401-k		ROTH 401-k	1,875
Total deductions	18,425	Total deductions	18,188
Net salary	22,553	Net salary	21,812

Table 4: Leslie's Budget: Comparison of Pre- and Post-Retirement Savings Deductions (\$)

Ultimately, the choice between a traditional and ROTH savings account boils down to

the following considerations.

- **Liquidity.** Do we need more money today? If so, a traditional IRA reduces taxes today, and provides more income.
- **Future wealth.** How large will our withdrawals be in the future? If large, then traditional 401-k withdrawals may be subject to high taxes, leaving us less money in retirement.
- **Future tax rates.** Do we expect future tax rates to go up or down? If they go up significantly, this could really eat into traditional 401-k withdrawals, which are taxed as ordinary income regardless of whether the growth came from capital gains, dividends, or interest income.

These are difficult questions to answer. But, they are important to consider when choosing between the options.

Traditional and ROTH 401-k's, 403-b's, IRAs are known as **defined contribution** plans. Most employer sponsored retirement plans in the U.S. are of this type, in which it's the employee's responsibility to contribute to the plan and invest the money. How much we get out of our defined contribution plan at retirement is determined entirely by how much we choose to contribute - up to legal limits - and our skill at investing the money. Contribute a lot of successfully invest the money, we get a lot in retirement. Contribute less, invest poorly, we get less. An important consideration to these plans is that many employers will match our contributions. This is important! Matches are basically additional money our employer gives us so we need to should try to maximize the match according to the rules of our plan.

An alternative is a **defined benefit** plan or **pension** which is managed by the employer. Retirement payments are determined by the employer using a formula based on how long we've worked with the company, our age, and our salary. Thus, retirement payments are the responsibility of the employer. The nice thing about pensions is that the employer shoulder's the risk. We don't have to worry about contributing or investing. That said, we also don't have a choice about contributing or investing strategy, and if we switch jobs it could have a significant impact on our pension benefits.

Bottom line: We need to pay close attention to our employer benefits - retirement, health, etc. They are an important component of any compensation package.

A.5 Expenses

Now that we have our income, we can consider our expense - where and how much of that income we spend. We'll organize expenses into two groups.

1. **Nondiscretionary expenses.** Expenses we cannot avoid.
2. **Discretionary expenses.** Expenses we can avoid.

The line between the two groups is blurry and varies from person to person. Nonetheless, it provides a useful taxonomy and forces us to decide on what we can and can't live without.

1.5.1 Nondiscretionary Expenses

Table 5 presents all of Leslie's nondiscretionary expenses and the correspond values for January and March. Unlike her income, her expenses are more volatile. Many of the expenses occur infrequently. For example, estimated taxes are payable every quarter.³ Her life insurance premium is paid in full each November. Clothes are usually purchased in early Spring and Fall, and gifts tend to be purchased around birthdays and holidays.

Table 5 highlights the importance of a budget. Leslie's expenses vary dramatically throughout the year. In the spreadsheet accompanying this appendix, her non-discretionary expenses range from a low of \$8,114 in June to a high of \$33,614 in December. The only way to ensure she has enough money throughout the year is by having a budget and planning accordingly.

Additionally, while all of the expenses in table 5 are nondiscretionary, that doesn't mean they can't be reduced. For example, her family can shop at a less expensive grocery store, use the car less, purchase less expensive clothes and gifts, turn down the heater in the winter and air conditioning in the summer, stop the apartment cleaning service they use, etc. There is flexibility within some of these categories to spend less or more.

1.5.2 Discretionary Expenses

Table 6 details Leslie's discretionary expenses for January and March. These are the expenses Leslie can do without if necessary. Clear from the table, discretionary expenses vary even more than non-discretionary expenses.

³Her income and choice of tax deductions means she has to make quarterly estimated tax payments during the year to avoid owing the tax authority too much at the end of the year.

1.5.3 Credit Cards

Noticeably absent from your list of expenses is a credit card(s). Credit card bills are *not* expenses. They are loans for expenses. So, our budget should never have a line item corresponding to “credit card.” We need to look at the list of charges on our credit card and classify those into an appropriate category.

A.6 Does the Budget Make Sense?

Once set up, we have to ask whether the budget makes sense. Are we consistently spending more than we make? That strategy won't last long. Look at the budget and make sure it is realistic and sustainable. Consider whether each line item is an honest and accurate representation of what we expect to earn/spend. It's ok to run a deficit - expenses greater than income - in a month or two if we're able to cover it with money in our checking or savings account.

Table 7 presents an aggregated version of Leslie's budget so it fits on the page. A couple of comments:

1. Leslie expects to never run a deficit; her income is always greater than her expenses. Thus, Leslie is able to save money every month towards retirement or other unexpected expenses (e.g., medical, extra nights out to dinner, etc.)
2. There is a lot of variation in how much Leslie saves each month. In December, she has over \$54,000 dollars she can save, while in July she has \$2,075. This is neither good nor bad, just a consequence of her income and expenses through the year.

A.7 Following the Budget

Every month, Leslie needs to compare her actual income and expenses to her estimated income and expenses from her budget. Most every estimate will be wrong in that her actual expenses, and perhaps even some of her income, will be different from her budget's estimates. That's ok. She needs to understand where her money is *actually* going in contrast to where she thinks it's going. Additionally, understanding why she spent (or earned) more or less than what she thought she would enables her to respond by spending less or more in the future.

This adaptability is critical to avoid falling into a financial hole out of which it is impossible to climb out.

Table 8 presents Leslie's budgeted and actual income and expenses for March. We've aggregated some of line items so the table will fit on one page. Starting at the top, Leslie's income for March is \$3,000 below her expected income because she didn't receive a rent check. Why? If it was because her tenant failed to pay because of financial problems, this is a signal that she may not be able to count on future rental income and should adjust her spending appropriately. If instead, she didn't check the mail or forgot to deposit the check, then this is less of a concern and the \$3,000 should show up next month.

Leslie went \$1,805 over her budgeted nondiscretionary expenses - about 13.8%. This was due mainly to excess spending on groceries, clothes, and gifts. Why did she go over budget on these items? Were her budget estimates too low? Did she and her husband get carried away? Was there an unplanned event (visitor, party, etc.) that added to the usual costs? If so, should she be planning for these unplanned events because they are relatively frequent, even though somewhat unpredictable?

Leslie was under her budget for discretionary expenses by \$320. Even though close to her budget, it's instructive to look at the breakdown of expenses. It turns out that airfare was \$500 more expensive than she had thought, and because of weather, they didn't do as many activities as they had planned so saved \$400 on planned activities. When she and her family returned home, they had little desire to go out and saved \$420 on entertainment. Her time at the spa negated the need for a salon visit; hence, the savings on health & wellness.

There are several potential lessons from this comparison. Perhaps she needs to budget more money for airfare, or book further in advance or less costly flights. Maybe she can budget for less entertainment and salon expense in months where they go on vacation.

	Jan	Mar
Reinvested earnings		
Estimate taxes	10,000	
Housing		
Rent	6,000	6,000
Utilities (Gas, electric, water)	150	150
Cleaning	200	200
Internet + streaming	170	170
Groceries	850	850
Phone	129	129
Insurance		
Renters	175	175
Umbrella	120	120
Life		
Car		
Gas	75	75
Insurance	145	145
Dog		
Food	100	100
Vet	200	
Medical (Dr., Rx, Glasses)		240
Clothes		
Parents		3,000
Kids		1,200
Gifts		
Family		500
Services		
Total	18,314	13,054

Table 5: Leslie's Budget: Nondiscretionary Expenses for January and March (\$)

	Jan	Mar
Vacation		
Airfare		
Hotel		3,500
Food		2,100
Activities		1,600
Kids camp		
Entertainment		
Dining out	700	700
Movies, museums,...	100	
Health & wellness		
Beauty	200	200
Gym	140	140
Total	1,140	8,240

Table 6: Leslie's Budget: Discretionary Expenses for January and March (\$)

Month	Jan Mar												
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	
Total Income	25,553	25,553	25,553	25,553	27,429	27,429	27,429	27,429	28,657	28,657	28,657	28,657	101,057
Nondiscret. Exp.	18,314	8,714	13,054	18,314	9,214	8,114	18,314	12,314	8,114	19,054	12,114	33,614	
Discret. Exp.	1,140	3,610	8,240	960	1,040	7,460	7,040	1,460	1,040	960	1,040	12,460	
Total Exp.	19,454	12,324	21,294	19,274	10,254	15,574	25,354	13,774	9,154	20,014	13,154	46,074	
Savings	6,099	13,229	4,259	6,279	17,175	11,855	2,075	13,655	19,503	8,643	15,503	54,983	

Table 7: Leslie's Budget: Discretionary Expenses for January and March (\$)

	Mar	
	Est	Act
<hr/>		
Income		
Net salary	22,553	22,553
Securities income		
Rental income	3000	0
Total income	25,553	22,553
<hr/>		
Nondiscretionary expenses		
Reinvested earnings		
Estimate taxes		
Housing	6,520	6,534
Groceries	850	1,100
Phone	129	160
Insurance	295	295
Car	220	205
Dog	100	460
Medical (Dr., Rx, Glasses)	240	320
Clothes	4,200	5,100
Gifts	500	685
Total nondiscretionary expenses	13,054	14,859
<hr/>		
Discretionary expenses		
Vacation	7,200	7,500
Kids camp		
Entertainment	700	280
Health & wellness	340	140
Total discretionary expenses	8,240	7,920
<hr/>		
Savings/Deficit	4,259	-226

Table 8: Leslie's Budget: January and March Estimated vs. Actual (\$)

Appendix B

Financial Accounting

This appendix is a detour into accounting for those that need it. While not required to understand anything in the text, the concepts discussed here are useful on their own. Financial information is stored and reported in financial statements (e.g., income statement, balance sheet, cash flow statement). We're not going to learn how to construct financial statements here. That's what accountants do. We're going to learn how to use the statements.

We'll use Microsoft as our illustrative vehicle. Microsoft operates on a June fiscal year end meaning the statements cover the period July 1, 2020 to June 30, 2021. Don't misinterpret the discussion here as only applying to large, publicly-traded, tech companies. Financial statements for all companies are broadly similar and their analysis proceeds along similar lines. What differs are the numbers and the story they tell.

B.1 Income Statement

The **income statement** goes by several names including **statement of operations**, **statement of income**, and **profit and loss statement** or **P&L**. It tells us the company's sales, expenses, and profits during a period, typically a quarter or a year. Income statements can also be constructed for divisions and even projects within a company. Microsoft's income statement is presented in [Table 1](#)

Because of flexibility provided to companies in how they present their financial statements, different companies will often have different row labels or **line items** in their income statements. While there is a consistent theme to all income statements, this inconsistency can make cross-company comparisons difficult. In response, financial analysts developed a common income statement format that consists of several key metrics and that can be used

(\$millions)	2021
Revenue:	
Product	\$71,074.0
Service and other	97,014.0
Total revenue	168,088.0
Cost of revenue:	
Product	18,219.0
Service and other	34,013.0
Total cost of revenue	52,232.0
Gross margin	115,856.0
Research and development	20,716.0
Sales and marketing	20,117.0
General and administrative	5,107.0
Impairment, integration, and restructuring	0.0
Operating income	69,916.0
Other income (expense), net	1,186.0
Income before income taxes	71,102.0
Provision for income taxes	9,831.0
Net income	\$61,271.0

Table 1: Microsoft 2021 Income Statement (\$mil)

to compare different companies, or the same company at different points in time, more easily. Figure 2 presents Microsoft's income statement in that format, along with some lingo and definitions. Numbers in parentheses correspond to negative numbers.

We'll walk through table 2 line by line, but before doing so, the basic mechanics of the income statement are as follows.

1. Start with sales, which correspond to all the products and services sold during the period.
2. Subtract an expense (negative expenses correspond to income), which correspond to all the costs the company incurred to generate the sales for the period.
3. Compute a measure of earnings (a.k.a., income, profits).
4. Repeat steps 1. through 3. until you run out of expenses.

For example, Gross profit equals Sales minus Cost of sales. EBITDA equals Gross profit minus SG&A. And so on.

(\$millions)	2021	Lingo/Definitions
Sales, net	\$168,088	a.k.a., Revenue, turnover, receipts
Cost of sales	40,546	a.k.a., COS, Cost of revenue (COR), Cost of goods sold (COGS)
Gross profit	127,542	
SG&A	45,940	Selling, general & administrative expenses. a.k.a. overhead, fixed costs
EBITDA	81,602	Earnings Before Interest, Taxes, Depreciation, and Amortization
Depreciation & amortization	11,686	
EBIT	69,916	Earning Before Interest and Taxes. a.k.a., operating earnings/income
Other expenses (income)	(1,186)	Income/expenses unassociated with core business
Pre-tax income	71,102	Income used to compute taxes owed
Taxes	9,831	Tax expense
Net income	\$61,271.0	a.k.a., earnings, profit (loss), the bottom line

Table 2: Microsoft 2021 Income Statement Reorganized (\$mil)

To draw an analogy to personal finance, each of us has a personal income statement. Each year (or month), we receive income from our employer from effectively selling our labor - just like a company sells goods and services. We then spend that money on rent, utilities, food, etc - just like companies spend money on rent, utilities, etc. Also like companies, we pay taxes. What remains after all those expenses are subtracted out of our income are our earnings that we can use as we please, e.g., save. There are some differences between a corporate and a personal P&L, as we'll see below, but the intuition is otherwise similar.

2.1.1 Sales

The income statement starts with **Sales**, otherwise known as **Revenue**, **Turnover**, **Receipts** or the “**top line**.” Net sales is gross sales less any returns, discounts, or refunds. For example, when we purchase a Surface tablet, Microsoft records a sale and gross sales increase. If we return the tablet 30 days later, the previously recorded sale doesn't disappear. Rather, an adjustment for the return is made and shows up in net sales.

Most firms employ **accrual basis** - as opposed to **cash basis** - accounting. This means that sales are **recognized** (a.k.a., **booked**, **recorded**) when they occur, *not* when money changes hands. This distinction is important for finance, which cares about when cash changes hands, not when transactions are recorded by accountants. To illustrate the potential disconnect between a sale and cash inflow, consider the following two transactions.

1. We buy a gaming console for \$1,000 with a credit card. As soon as the transaction is complete, Microsoft recognizes \$1,000 of revenue even though it won't receive the

money from the credit card company for a few days.

2. We buy a three-year, \$300 million contract for cloud computing services. Accountants recognize the revenue according to the percent of the contract that is completed each year. If our usage is uniform (equal) over the life of the contract, the accountants will book \$100 million of sales in each year of the contract. However, we may have negotiated an entirely different payment scheme with Microsoft, such as paying \$245 million up front at the start of the contract or \$340 million at the end of the contract. Figure B.1 illustrates this distinction with a timeline. Regardless of when we pay, Microsoft recognizes the revenue according to the usage of the service.




Figure B.1: Sales versus Cash Inflows

Both examples highlight an important point: sales do not always coincide with cash flowing into the company unless the sale was paid for with cash.

2.1.2 Operating Expenses

Expenses on the income statement are the costs of generating revenue. There are two types: (i) operating, and (ii) non-operating. The former corresponds to expenses associated with a company's core business - the activities primarily responsible for their current and future revenue. There are three categories of operating expenses, the first of which is **cost of sales (COS)**.

Also known as **cost of revenue (COR)** and **cost of goods sold (COGS)**, cost of sales captures the direct costs of selling goods or services. Cost of sales is sometimes referred to as **variable costs** because these costs tend to move closely with output - the *quantity* of goods or services sold. According to Microsoft's annual filing (10-K) with the Securities Exchange Commission (SEC), costs of sales include:

Cost of revenue includes: manufacturing and distribution costs for products sold and programs licensed; operating costs related to product support service centers

and product distribution centers; costs incurred to include software on PCs sold by original equipment manufacturers (“OEM”), to drive traffic to our websites, and to acquire online advertising space; costs incurred to support and maintain online products and services, including datacenter costs and royalties; warranty costs; inventory valuation adjustments; costs associated with the delivery of consulting services... (Microsoft 10-K, 2021, page 65)

The second category is **SG&A**; **selling, general, and administrative** expenses. While the name is suggestive, Microsoft states that the “G&A” in this line item includes:

General and administrative expenses include payroll, employee benefits, stock-based compensation expense, severance expense, and other headcount-related expenses associated with finance, legal, facilities, certain human resources and other administrative personnel, certain taxes, and legal and other administrative fees. (Microsoft 10-K, 2021, page 46)

Marketing and research and development (R&D) expenses are also part of SG&A, which are often referred to as **fixed costs** because of their relative insensitivity to changes in output. Utility bills, managerial salaries, and internet costs are less sensitive to a company’s output. That said, the delineation between cost of sales and SG&A is not sharp. Accountants have some discretion over where they classify different expenses.

From a finance perspective, the classification of an expense is less important than the nature - variable or fixed - of the expense. Variable expenses act as a **natural hedge** against sales risk. In other words, when sales decline, variable costs decline thereby easing pressure on earnings and cash flow. Fixed expenses do not decline when sales decline and therefore increase pressure on earnings and cash flow. The ratio of fixed to variable costs is referred to as **operating leverage**, which measures operational risk stemming from the company’s cost structure. The higher the operating leverage, the more volatile and risky the earnings.

As with sales, both cost of sales and SG&A need not align with when money is paid. Using our cloud computing service example above, the costs of the sale, such as salaries, are recorded when the corresponding revenue is recognized, not when the money is actually paid. As another example, if Microsoft pays \$200 today for materials used to construct a gaming console it sells six months later, the \$200 is only recognized as an expense when the sale occurs. This matching of expenses to sales is consistent with the revenue recognition discussed above and a central theme of accrual accounting.

The final category of operating expenses is **depreciation and amortization**.¹ When Microsoft buys an asset (e.g., robot, plant, real estate, company, patent), the money spent does not appear on the income statement. Rather, the cost of the asset shows up on another financial statement, the balance sheet, which we discuss below. After the purchase, the asset value is expensed via depreciation or amortization over its usable life. Each year *after* the purchase, a fraction of the asset's value shows up as an expense on the income statement, and the value of the asset on the balance sheet is reduced by the same amount.² Whether the company depreciates or amortizes the asset depends on whether the asset is tangible or not. If you can touch it (plant, property, equipment), you depreciate it. If you can't touch it (patent, licensing agreement, trademark, copyright, software), you amortize it.

For example, imagine Microsoft spends \$10 million to buy a robot with a five-year usable life. At the end of five years, the robot is anticipated to be worth \$1 million dollars in **salvage value** (i.e., what the asset is worth at the end of its usable life). To **straight-line depreciate** an asset, accountants would compute the periodic value loss like so

$$\text{Depreciation} = \frac{\text{Purchase price} - \text{Salvage value}}{\text{Length of usable life}} \quad (\text{B.1})$$

In our example, the annual depreciation is $(10 - 1)/5 = \$1.8$ million. In other words, after acquiring the robot, Microsoft will report \$1.8 million of depreciation expense on its income statement each year for the following five years.³

Depreciation and amortization are **non-cash expenses**. No money leaves the company when an asset is depreciated or amortized. In some sense, it's just accountants way of recognizing that assets tend to decline in value over time. Of course, some assets, such as real estate, can appreciate over time. And, others can quickly become obsolete as a result of technological innovation. So, its important not to confuse depreciation and amortization - accounting concepts - with actual declines in value - financial concepts. Sometimes they provide a reasonable approximation to wear and tear on assets, other times they do not.

¹To be precise, the asset must be part of the company's normal or core business operations in order for the corresponding depreciation and amortization expense to be considered a part of operating expenses.

²Asset must have a have a finite life to be depreciated. Land is an example of an infinitely lived asset that doesn't depreciate. Likewise, current assets on the balance sheet are not depreciated (see below).

³In practice, accountants in the U.S. compute depreciation on an accelerated basis using the Modified Accelerate Cost Recovery System (MACRS). This approach front-loads depreciation expense relative to the straight-line approach.

2.1.3 Non-Operating Expenses

Other and Taxes represent the non-operating expenses on the income statement. The Other category includes all of the expenses - and income - a company generates from operations that are not central to its business. For Microsoft, this includes mostly financial investments as detailed in table 3.

(\$millions)	2021
Interest and dividends income	\$2,131
Interest expense	(2,346)
Net recognized gains on investments	1,232
Net gains (losses) on derivatives	17
Net losses on foreign currency remeasurements	54
Other, net	98
Total	\$1,186

Table 3: Other Income (Expenses) Microsoft 2021 Income Statement Reorganized (\$mil)

Microsoft invests in bonds and stocks that generate interest and dividend income. They also pay interest on what they've borrowed (interest expense). In addition, they “use derivative instruments to: manage risks related to foreign currencies, equity prices, interest rates, and credit; enhance investment returns; and facilitate portfolio diversification.” The Other category can also include nonfinancial expenses and income so long as they are unrelated to the core business of the company, such as the cost of disposing of equipment.

Taxes correspond to taxes owed on income, as opposed to taxes owed on other assets like property. Taxes are computed using a firm's taxable income indicated by the Pre-tax income line. The larger the pre-tax income, the larger the tax bill. So, companies have an incentive to reduce their taxable income to avoid paying taxes. Though, they trade this benefit against the cost of reporting lower earnings.

The tax expense reported on the income statement is rarely the tax paid by the company, and not just because it contains a subset of taxes. There are different accounting rules required by tax authorities, like the Internal Revenue Service (IRS) in the U.S., and the accounting principles used to prepare financial statements. Trying to uncover the true tax bill from financial statements is difficult if not impossible. So, the tax expense item needs to be taken with a grain of salt.

2.1.4 Earnings (a.k.a., Profits)

After each expense comes a measure of earnings or profit.

- **Gross profit** tells us how much the company earns after direct (or variable) costs are removed from sales.
- **Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA)** is one measure of **operating income** or **operating earnings**, that is how much money the company is making after deducting all cash operating expenses. (Remember that depreciation and amortization are *non-cash* expenses.) It is also a measure of the cash earnings available for the firms creditors and shareholders. EBITDA is a popular measure because it is unaffected by items that are unrelated to the core operations including:
 - Depreciation and amortization are determined largely by accounting rules and the nature of the business' assets.
 - Other expenses and income are by definition unrelated to the core operations of the company and are often largely influenced by the financial policy (choice of debt and equity) of the company.
 - Taxes are determined by the government.

Because EBITDA is unaffected by these factors, it is useful for comparing the operating performance across different companies. It is also useful for measuring the income available to repay debt and, as such, frequently appears in loan contract provisions known as covenants.

- When people refer to operating income or operating earnings, they are typically talking about **Earnings Before Interest and Taxes (EBIT)** because management often view depreciation and amortization as corresponding to the real economic costs of utilizing assets. Like EBITDA, EBIT measures the income available to creditors and equity holders, and is unaffected by the firm's financial policy, non-core business expenses/income, and tax policy.
- When people talk about **earnings**, they are talking about **net income** or “**the bottom line**” of the income statement. Net income measures the after-tax earnings available for stock holders should the firm wish to pay a dividend or buy back shares. These earnings can also be retained by the firm for future use.

Hopefully clear from their descriptions, each measure of earnings contains somewhat different information about the company. One point to never lose sight of is that none of the earnings measures exactly coincide with cash moving in or out of a company, though they can still be informative.

B.2 Balance Sheet

The balance sheet tells us what the company owns (**assets**) and what they owe (**liabilities and shareholders equity**). The balance sheet is a snapshot at a point in time, as opposed to a recording over a period of time like the income statement. Figure 4 presents Microsoft's balance sheet as of June 30, 2021. It is often presented in one long column, like the income statement, with assets on top and liabilities & shareholders equity on the bottom. I've split it into two sides to conserve space. The left side of the balance sheet details the assets the company owns. The right side of the balance sheet details the liabilities and shareholders equity the company owes.

Assets	2021	Liabilities & Shareholders Equity	2021
Cash	\$130,334	Accounts payable	\$15,163
Accounts receivable, net	38,043	Accrued compensation	10,057
Inventories	2,636	Unearned revenue	41,525
Other	13,393	Debt and other	21,912
Current assets	184,406	Current liabilities	88,657
Net PP&E	59,715	Long-term debt	59,703
Goodwill & intangibles	57,511	Other	43,431
Financial and other	32,417	Total liabilities	191,791
Total assets	\$333,779	Common stock & paid-in capital	83,111
		Retained earnings	58,877
		Shareholders equity	141,988
		Total liabilities and equity	\$333,779

Table 4: Microsoft 2021 Balance Sheet (\$mil)

Like the P&L, each of us has a personal balance sheet. On the asset side, there's the stuff we own: cash in a checking or savings account, money owed to us by our employer for the work we've already provided, a house, stock, and bond investments. On the liability side is how we paid for all those assets: a credit card balance, a personal loan, a mortgage. And, finally, there's our equity or **net worth**, our money.

2.2.1 Assets

Companies own assets with the hope that they will generate future profits. The different assets contribute in different ways towards that goal. Assets that don't contribute are deemed **excess assets** and are often sold.

Current Assets

Current assets are those that can be converted into cash within one year, so-called **liquid** assets. The most liquid asset of course is **cash** itself, and Microsoft has a lot of it, \$136 billion! To be precise, this isn't all just cash sitting in a bank. It consists primarily of **short-term investments** and **marketable securities**, financial assets that can be quickly and easily converted to cash. Examples include Treasury bills, commercial paper, promissory notes, and money market accounts. Cash is used by companies to run day-to-day operations, make strategic investments, and as a buffer against bad times and earnings shortfalls.

Accounts receivable, or **receivables**, represent money owed to Microsoft by its customers. When we buy something from a company and pay by means other than cash (e.g., credit), there is a disconnect between when the company recognizes the sale - at the time of the sale - and when they receive the cash - sometime later. As long as the company anticipates receiving the money from the sale within a year, that money gets recorded in accounts receivable. Microsoft's customers owe them \$38 billion that the company expects to receive within a year and in most instances a lot sooner.

For credit sales to consumers who pay by credit card, the credit card company will pay the selling company within one to three days, though larger transactions can take longer. For credit sales to businesses, the terms of the credit can vary widely but 30, 60, and 90 day terms are common.

Inventories consist of the materials and products a company has yet to sell. There are three stages of inventory: (i) raw materials, (ii) work-in-progress (WIP), and (iii) finished goods. The bulk of Microsoft's revenues comes from software and services, which is why their inventory (\$2.6 billion) is small relative to their other assets and their sales (\$168 billion).

Long-Term Assets

Long-term assets are not expected to be converted into cash within one year. **Tangible assets**, i.e., stuff you can touch, consist of **PP&E** or **Property, Plant, and Equipment**. This account comes in two flavors, gross and net. Let's see how they work with an example.

Imagine Microsoft buys the robot described earlier at the end of 2020. Also assume it's the only asset Microsoft owns, and they don't buy anymore assets. **Gross PP&E** and net PP&E for 2020 increase by \$10 million, the purchase price. One year later, at the end of 2021, the robot has experienced \$1.8 million of depreciation, which shows up on the 2021 income statement. The 2021 balance sheet shows the same gross PP&E, assuming nothing else was purchased during the year, but net PP&E will decline by \$1.8 million to \$8.2 million to reflect the depreciation. In 2022, gross PP&E still shows \$10, but net PP&E declines by another \$1.8 million to \$6.4 million. And so on. So, gross PP&E is a running total of the long-term tangible assets the company has purchased, **net PP&E** is that total less **accumulated depreciation**.

Intangible assets are the stuff you can't touch, like patents, trademarks, copyrights, software, and so on. **Goodwill** is another intangible asset, and it has nothing to do with charity. When companies buy assets they sometimes pay a **premium**, or more than the fair market value of the asset. This situation often arises in the acquisition of another company. In 2016, Microsoft purchased LinkedIn for \$26.2 billion, but LinkedIn's assets were worth a fraction of that, approximately \$18 billion. The difference of $26.2 - 18.0 = \$8.2$ billion is recorded as Goodwill.

Goodwill is typically not amortized. Instead, it is tested each year for impairment. An asset is deemed impaired if its current market value is less than the **carrying value**, or the value on the balance sheet. When an impairment occurs, the value of the asset on the balance sheet is reduced, or **written down**, and an expense is shown in the Other expenses line of the income statement. In 2015, Microsoft had to take an **impairment charge** of \$7.6 billion of assets associated with its acquisition of Nokia in 2013.

Total assets is the sum of current and long-term assets, $184,406 + 59,715 + 57,511 + 32,147 = \$333,779$. (Long-term assets are not sub-totaled.)

2.2.2 Liabilities & Shareholders Equity

Companies can pay for their assets in a variety of ways, all of which are detailed on the right side of the balance sheet.

Current Liabilities

Money owed within one year is referred to as a **current liability**. The most common current liability appearing on balance sheets is **accounts payable** or **payables**. This account shows

how much companies owe to their suppliers and is the counterpart to accounts receivable. Much like customers purchase goods and services on credit, companies do the same.

Accrued compensation and **unearned revenue** are somewhat less common but worth describing. The former account corresponds to money earned by our employees, but not yet paid. The latter account corresponds to money received from a customer for a service not yet provided or a good not yet delivered.

The other current liability common to most balance sheets is **Debt** or, more precisely, **short-term debt** and **long-term debt due**. Short-term debt contains loans with maturities less than one year. Long-term debt due contains loans with maturities greater than one year but needing to be repaid within one year. These two accounts are distinct from other liabilities in that they require the company to pay a market rate of interest. There is no interest expense on the other accounts we mentioned.⁴

The difference between current assets and current liabilities is called **net working capital** or simply **working capital**. Working capital measures a company's **liquidity**, net of what it owes in the short-term, from its operations. Companies with lots of working capital have lots of assets that are or will shortly become cash relative to what they owe over the next year. Microsoft's working capital is $184,406 - 88,657 = \$95,749$.

Long-term Liabilities

Liabilities not classified as current correspond to money owed after one year. **Long-term debt** is debt with maturities greater than one year and is a common account on many balance sheets. Like its current liability counterpart, what distinguishes long-term debt from most other long-term liabilities is that debt earns a market rate of interest. (Leases are another example in which a market rate of interest is earned, but this is a relatively small number for Microsoft that is aggregated in the Other line item.) Examples of other long-term liabilities are:

- income taxes,
- unearned revenue, and
- deferred income taxes.

⁴Accounts payable and receivables, sometimes referred to as **trade credit**, offer terms that look like interest. For example, if a company pays its bill within 10 days, it is offered a 2% discount on its bill. If they don't pay within 30 days, they are deemed late in their payment and potentially subject to steep penalties. These terms are referred to as "2/10 net 30."

The key characteristic of these accounts from an accounting perspective is that they correspond to monies owed outside of one year from the date of the balance sheet.

Shareholders Equity

Companies can buy assets by borrowing money, recorded as liabilities. Or, they can buy them with their owners' money, recorded as shareholders equity. This situation is completely analogous to buying a house. The home loan (mortgage) is a liability. The money put down by the owner is equity.

Perhaps more important to understand than what is shareholders equity on the balance sheet, is to understand what it isn't. Shareholders equity on the balance sheet, sometimes referred to as **book equity**, is not the same as a company's market capitalization or market value of equity. In fact, book equity often bears little resemblance to market capitalization. Microsoft has a book equity of \$141,988 billion, but its market cap is \$2.0 *trillion* as of June 30, 2021 . Many companies have negative book equity, which cannot occur for market capitalizations because of limited liability.⁵

The primary role of shareholders equity on the balance sheet is to ensure that the balance sheet "balances." That is,

$$\text{Total assets} = \text{Total liabilities} + \text{Shareholders equity} \quad (\text{B.2})$$

This is an **accounting identity** and it holds for every balance sheet you will see. That said, let's understand the two equity accounts on Microsoft's.

U.S. states require common stock to have a **par** or **stated** value, much like the par or face value of a bond. This doesn't mean much for investors who are only concerned with the market value. But, it's relevant for accounting. Imagine a company that issues stock with a par value of \$1 per share. An investor purchases it for \$10. The company recognizes this transaction on the balance sheet by increasing the common stock account by \$1, and the paid-in-capital account by \$9. Figure 4 shows these two accounts combined for Microsoft.

Retained earnings equal the company's cumulative net income from the income statement minus the cumulative amount of dividends paid over the same period. The cumulation begins at the date of the company's incorporation and ends at the current balance sheet date. For example, when the company reports net income of \$10, retained earnings increase by \$10. A \$10 loss reduces retained earnings by \$10, just like a \$10 dividend.

⁵Limited liability prevents creditors from going after the personal assets of shareholders. So, the worst that can happen to a shareholder is that their stock becomes worthless.

B.3 Cash Flow Statement

The last financial statement we'll look at is the cash flow statement, whose purpose is to show how much money is flowing through the company and how it flows. This sounds like what the income statement should do but remember two features of the income statement. First, it relies on the accrual method of accounting and records transactions, not the movement of cash. Second, it records income and expenses, not investments like the purchase of inventory or equipment. We can view the cash flow statement as undoing accrual accounting so we can see the actual dollars and cents moving in and out of a company. Figure 5 presents Microsoft's 2020 cash flow statement.

	2021
Operations:	
Net income	\$61,271
Depreciation, amortization, and other	11,686
Accounts receivable	(6,481)
Inventories	(737)
Accounts payable	2,798
Stock-based compensation expense	6,118
Other	2,085
Net cash from operations	<u>76,740</u>
Investing:	
Additions to property and equipment	(20,622)
Purchases of intangible and other assets	(8,909)
Purchases of investments	(62,924)
Maturities of investments	51,792
Sales of investments	14,008
Other, net	(922)
Net cash used in investing	<u>(27,577)</u>
Financing:	
Repayments of debt	(3,750)
Common stock issued	1,693
Common stock repurchased	(27,385)
Common stock cash dividends paid	(16,521)
Other, net	(2,552)
Net cash from (used in) financing	<u>(48,515)</u>
Net change in cash	<u>\$648</u>

Table 5: Microsoft 2021 Statement of Cash Flows (\$mil)

There are three sections to the statement showing cash flow from different parts of the company: operations, investing, and financing.

2.3.1 Operations

To get the cash flows from operations, we start with net income from the income statement. We then add back the non-cash expenses (depreciation and amortization) and changes in working capital (e.g., accounts receivable, inventory, accounts payable).

Consider depreciation and amortization. We subtracted this from sales on the income statement. But, no money actually left the company as a result of depreciation and amortization. We only do it for tax purposes. So, we add back depreciation and amortization to get closer to the actual money flowing in or out of the company. Stock-based compensation is another non-cash expense. The company pays its employees with both cash and stock. The stock is considered an expense that is deducted from sales on the income statement. (It's in SG&A.) However, no cash leaves the company as a result of this stock based compensation. So, we add it back.

Similarly, when we book a sale paid on credit, we don't receive any money. As a result, accounts receivable increases by the amount of the sale. When we receive the money later, accounts receivable goes down and our cash balance goes up. Microsoft's accounts receivables increased by \$6.481 billion from 2020 to 2021. This increase could have happened for a variety of reasons including increased sales, longer payment terms for Microsoft's customers, or changes in Microsoft's customer behavior.⁶ Regardless, the net effect of an increase in accounts receivable is that Microsoft is *not* getting money. Thus, we subtract this change from net income.

Inventories went up by \$737 million, meaning Microsoft increased the product in its warehouses. This increase means that even more cash was invested in inventory. So, we subtract the increase in inventory from net income.

More generally, changes in balance sheet accounts have the following effects:

- Increases in assets coincide with reductions in cash; decreases in assets coincide with increases in cash.
- Increases in liabilities coincide with increases in cash; decreases in liabilities coincide with decreases in cash.

⁶Accounts receivable and accounts payable are typically reported on a "net" basis, where net refers to net of bad debts.

To illustrate the second point, accounts payable increased by \$2.798 billion between 2020 and 2021. This could have happened because Microsoft needed to purchase more products and services from their suppliers to support greater sales, or they received longer payment terms from their suppliers. Regardless of the reason for the increase, Microsoft is retaining more cash by not paying their suppliers. Eventually, they will of course have to make these payments, just like their customers will have to pay them. However, until that time, Microsoft keeps the cash and its cash flow increases.

The net results of their operations was to generate \$76.740 billion of cash in fiscal 2021.

2.3.2 Investment

The investment section details the money made and spent on all types of investments. Microsoft spent \$20.622 billion on tangible investments in property and equipment, and another \$8.909 billion on intangible investments. Financial investments play a large role at Microsoft, as suggested purchases, sales, and maturation of their investments. In total, Microsoft spent \$27.577 billion on investing activities.

2.3.3 Financing

The last section details the cash flows from Microsoft's financing activities. This may sound similar to the financial investments discussed just above. But, financing refers to how they fund those, and other, investments as well as the broader operations of the company.

In 2021, Microsoft did not issue any new debt, though they repaid \$3.750 billion. They issued and repurchased stock, in addition to paying their shareholders dividends. Net of stock issuances, Microsoft distributed \$42.213 billion to shareholders.

In total, Microsoft spent \$48,515 billion on financing activities.

2.3.4 Summary

The net effect of the operations, investment, and financing is that Microsoft increased its cash holdings by \$648 million in 2021. Taking a step back, we should really appreciate the statement of cash flows now. It is telling us the *actual* money moving in and out of the company and through which channels.

B.4 Application: Financial Statement Analysis

Now that we understand how to read and interpret the big three financial statements, let's think about how to analyze them. To do so, we'll focus on different groups or categories of **key performance indicators (KPIs)** that provide insight into companies behaviors.

In addition, we'll compute the KPIs for Microsoft in 2019, 2020, and 2021. For comparison, we'll do the same for two of its closest competitors, Alphabet (ticker=GOOG) and Amazon (ticker=AMZN). These figures will provide both historical and competitive context to help interpret the most recent KPIs.

A few caveats before proceeding. First, as noted earlier, Microsoft is on a June fiscal year, meaning its financial statements are as of June 30 each year. Alphabet and Amazon are both on December fiscal year, meaning their financial statements are as of December 31 each year. Thus, there is a six month gap in the timing of Microsoft's financials and its competitors. We could resolve this by focusing on Microsoft's trailing twelve month (TTM) financials as of December 31, or Alphabet and Amazon's TTM financials as of June 30. For our purposes, the additional precision is unnecessary.

Second, KPIs are related to one another. Changing one often changes others. Therefore, we can't focus on one number. We need to look at the story *all* of the numbers tell. Finally, financial statement analysis often doesn't provide answers. Instead, it directs us to the questions that must to be asked.

2.4.1 Performance

Accounting rates of return or **accounting returns** measure the financial performance of a company over an historical period and can be useful diagnostics. The accounting adjective is important because it distinguishes accounting performance from market performance (e.g., stock or bond returns). Market performance tells us what investors earn, and this need not align with accounting returns.

1. Return on assets (ROA).

$$ROA_t = \frac{\text{Net income}_t}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{61,271}{(301,311 + 333,779)/2} = 0.193 \quad (\text{B.3})$$

Assets is total assets from the balance sheet. Because earnings are produced over a period, one year in this case, the assets responsible for producing those earnings are

a combination of those we started with and those we ended with. Thus, we use an average of the start ($t - 1$) and end ($t - 1$) of period assets in the denominator.

Despite its common usage, this is *not* how ROA should be measured. Assets generate earnings that are enjoyed by *all* of a firm's claimants - creditors and shareholders. Net income only measures the earnings available to shareholders. A more sensible definition for ROA uses a measure of after-tax operating income that is available to both creditors and shareholders.

$$ROA_t = \frac{EBIT_t \times (1 - \tau)}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{69,916 \times (1 - 0.138)}{(301,311 + 333,779)/2} = 0.190 \quad (\text{B.4})$$

We've assumed Microsoft's effective tax rate, τ , is 13.8%.

ROA tells us how much after-tax income is generated per dollar of assets; the more, the better *all else equal*. The caveat is important because one way to increase returns as we've discussed is to take more risk! So, holding fixed the business risk, more ROA is better.

Microsoft generated \$0.19 of after-tax operating income per dollar of assets in fiscal 2021. In 2017, Microsoft's ROA was 9.4% and has increased each year since, suggesting that Microsoft is getting more value from its assets each year.

2. Return on investment (ROI). Also known as **Return on Invested Capital (ROIC).**

$$\begin{aligned} ROI_t &= \frac{EBIT_t \times (1 - \tau)}{((\text{Debt}_{t-1} + \text{Equity}_{t-1}) + (\text{Debt}_t + \text{Equity}_t)) / 2} & (\text{B.5}) \\ &= \frac{69,916 \times (1 - 0.138)}{((70,988 + 67775) + (118,304 + 141,988))/2} \\ &= 0.302 \end{aligned}$$

where Debt and Equity are the book values from the balance sheet. Like ROA, larger ROI is better - holding fixed the risk of the business - because it indicates that investors' money is generating even more money. For each dollar that creditors and shareholders invested in Microsoft as of 2021, Microsoft generated \$0.30.

You might think: "Wait. Investors only get \$0.30 for each dollar they invest? That sucks!" But, this is just one year. As emphasized in the text, investors are entitled to many years of earnings (cash flows to be more precise). So, that \$1 invested will lead to much more income over time.

3. Return on equity (ROE).

$$ROE_t = \frac{\text{Net income}_t}{(\text{Equity}_{t-1} + \text{Equity}_t)/2} = \frac{61,271}{(118,340 + 141,988)/2} = 0.471 \quad (\text{B.6})$$

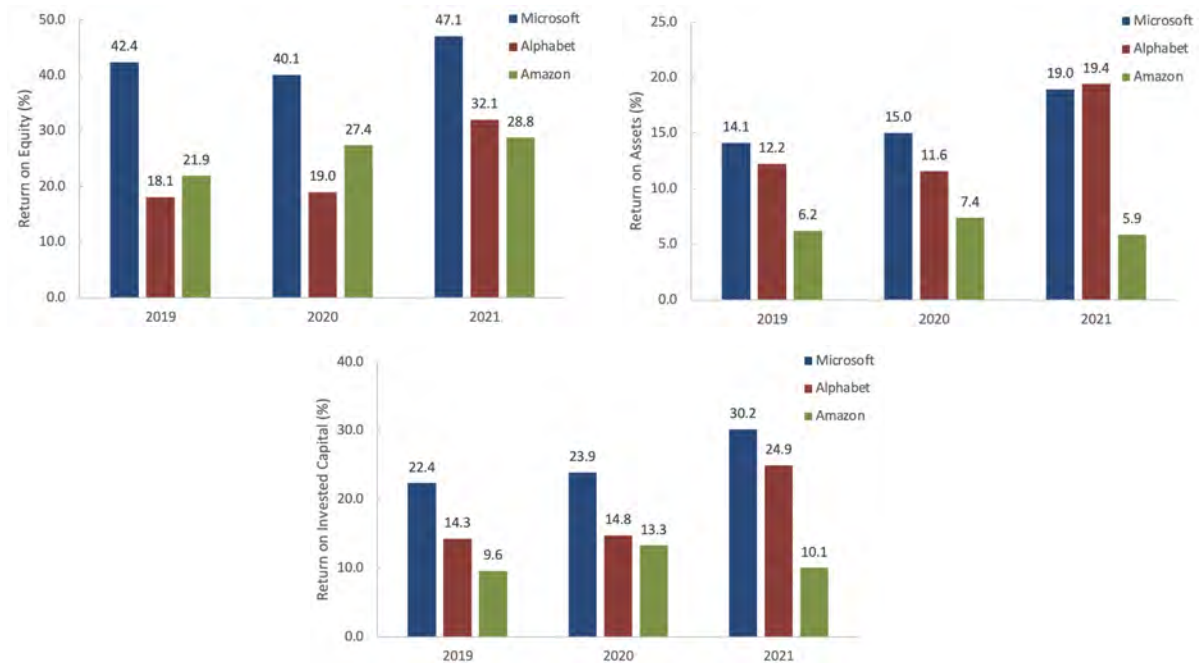


Figure B.2: Performance KPIs

ROE measures the income to shareholders per dollar they invest. Microsoft generated \$0.47 per dollar of equity investment in fiscal 2021. Again, larger ROE is better holding risk fixed.

An important caveat is that, unlike ROA and ROI, ROE is affected by a firm's capital structure because it is computed using net income which varies with interest expense. Firms with relatively more debt will, all else equal, tend to have higher ROEs simply because they are financially riskier even if they have similar business risk.

Figure B.2 presents current and historical performance KPIs for Microsoft and two of its closest competitors - Alphabet (ticker=GOOG) and Amazon (ticker=AMZN). The figure shows that Microsoft's performance in 2021 has been impressive by both historical and competitive standards. With the exception of 2020's ROE, all three return series have shown improvement over time. Also note the differences in the levels of the three measures. ROEs are significantly larger than ROAs and ROICs. This difference is due in part to the larger denominators in ROA and ROIC. It is also due to ROA and ROIC blending the returns to shareholders and creditors, the latter of which will be lower because of the lower risk faced by creditors.

Be careful not to misinterpret the cross-company comparisons as necessarily meaning that Microsoft's shareholders (or creditors) have been better off than those of Alphabet or

Amazon over this period. Differences in these performance measures - indeed *all* KPIs - largely reflect differences in business model even within the same industry.

Company	2019	2020	2021
Microsoft	38.0%	53.8%	34.4%
Alphabet	29.1%	31.0%	65.2%
Amazon	23.0%	76.3%	2.4%

Table 6: Annual Stock Returns

Table 6 shows the annual stock returns for each company over the same period. The difference between **accounting returns** and **market returns** is apparent. For example, despite a ROE that was almost 13% lower than Microsoft's in 2020, Amazon's stock return was 23% higher. Likewise, Amazon's ROE was at its highest in 2021. Yet, 2021 was its lowest stock return. While the correlation between accounting and market returns is generally positive, they do not measure the same thing. Accounting returns reflect what *has happened*. Market returns reflect what the market thinks *will happen*.

2.4.2 Profitability

Figure 7 presents Microsoft's income statements for the last three years in which every value is divided by sales. Expenses divided by sales are referred to as **expense ratios** and measure how much is spent on a particular expense per dollar of sales. Earnings divided by sales are referred to as **profit margins** or **margins** and measure the earnings per dollar of sales. Margins measure the profitability of a company.

Lower expense ratios and higher margins are preferable. But, as with all decisions, there are tradeoffs. Reducing expenses today may come at the cost of lower revenue growth tomorrow. Perhaps more interesting is that the expense ratios paint a picture of the cost structure of a company.

For Microsoft, the bulk of their expenses come from SG&A and cost of sales. This is not surprising as Microsoft is not a physical capital intensive (e.g., plants, equipment, etc.) business. Depreciation and amortization is a relatively small component of their costs. Rather, Microsoft makes software and provides services. Closer inspection of their 10-k reveals that most of the SG&A is engineering salaries. Therefore, cost improvements have to come from reductions in overhead (SG&A) and then costs of sales.

	2019	2020	2021	Lingo
Sales	100.0	100.0	100.0	
Cost of sales	24.8	23.3	24.1	COS expense ratio
Gross profit	75.2	76.7	75.9	Gross or contribution margin
SG&A	31.8	30.8	27.3	SG&A or overhead expense ratio
EBITDA	43.4	46.0	48.5	Operating margin
Depreciation & amortization	9.3	8.9	7.0	D&A expense ratio
EBIT	34.1	37.0	41.6	Operating margin
Other expenses (income)	(0.6)	(0.1)	(0.7)	
Pre-tax income	34.7	37.1	42.3	
Taxes	3.5	6.1	5.8	
Net income	31.2	31.0	36.5	Net or profit margin

Table 7: Microsoft 2020 Income Statement Reorganized (%)

Over time, Microsoft's cost structure has been relatively stable. However, Microsoft has shown consistent improvement in reducing their expense ratios. As a result, gross margin and operating margins (EBITDA and EBIT) have been increasing over time.

Microsoft's margins show a profitable company that is able to keep a significant fraction of their sales for their investors. Further, gross and operating margins have been increasing over the last three years, consistent with the decline in operating expense ratios. Figure B.3 shows operating and net margins for Microsoft, Alphabet, and Google from 2019 to 2021. The figures highlight how profitable Microsoft has been relative to its competitors.




Figure B.3: Profit Margins

The expense ratio and margin trends show that Microsoft appears to be on an impressive trajectory. Before drawing any conclusions, it is important to understand what is happening with sales growth. It could be that Microsoft is simply slashing expenses to boost margins and, in the process, sacrificing future growth. For example, Microsoft could fire half of our company's employees to drive down costs, but then there's not enough people around

to generate future sales! As a result, margins may go up in the short-run, but we'd be destroying value, which takes into account short- and long-run performance.

Figure B.4 presents revenue growth estimates. All three companies recent top-line growth rates are impressive in light of their scale. 2021 sales for Microsoft, Alphabet and Amazon are \$168, \$258, and \$470 billion, respectively. However, while Microsoft's sales growth is less impressive than its competitors, its greater profitability has led to greater returns for its investors, as shown above (figure B.2).




Figure B.4: Revenue Growth

Microsoft's ability to increase its sales while also increasing its profit margins has led to the second most valuable company on the planet as of April 2022 with a market capitalization of over \$2 trillion.

2.4.3 Asset Efficiency

Asset turnover or **asset utilization** ratios measure how efficiently firms utilize their assets to generate sales. We compute asset turnover ratios by taking the ratio of sales to average assets. For example, Microsoft's total asset turnover ratio of 2021 is

$$\text{Asset turnover}_t = \frac{\text{Sales}_t}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{168,088}{(301,311 + 333,779)/2} = 0.5$$

For each dollar of assets, Microsoft generates \$0.50 of sales.

In theory, we can compute turnover ratios for any asset. However, it really only makes sense to compute them for assets responsible for generating sales. For example, a deferred tax asset turnover ratio makes little sense.

Figure B.5 presents total asset and net PP&E turnover ratios for Microsoft and its competitors. We see that Microsoft's total asset turnover ratios have been stable during the 2019 to 2021 period. Coupled with its revenue growth over this period, this pattern in utilization ratios tells us that Microsoft has grown its assets at a similar rate to which it has grown sales.

This pattern is in contrast to its physical or tangible capital, as measure by the ratio of sales to average net PP&E. The declining pattern suggests its physical capital is growing more quickly than its sales and, as a result, Microsoft is getting less bang for the buck out of that capital. Is this bad? Not necessarily. Microsoft has grown its net PP&E at a rate of 28% per year since 2019. This corresponds to investment in *future* growth. That growth may take some time to reach top line revenues. The question this figure raises is: How long? At what point do those net PP&E utilization rates turn back up or level out?

It's also interesting to note the difference in utilization ratios between the three companies. Amazon has significantly higher total asset utilization ratios, while Alphabet has uniformly lower ones. Does this mean Alphabet is operating inefficiently compared to its competitors? Not necessarily. Remember, these differences reflect differences in business models to a large degree. Microsoft and Amazon have significant retail arms, whereas Alphabet is almost exclusively software and services.




Figure B.5: Asset Turnover/Utilization Ratios

2.4.4 Liquidity

Liquidity in a corporate context refers to company's assets that are easily converted to cash. Liquidity is important for companies because without it, they cannot cover their immediate

costs (e.g., salaries, utility bills, rent) and therefore cannot operate. Two common measures of **liquidity**, or the amount of assets easily converted to cash relative to short-term liabilities, are the following.

1. **Current ratio.**

$$\text{Current ratio}_t = \frac{\text{Current assets}_t}{\text{Current liabilities}_t} = \frac{184,406}{88,657} = 2.1 \quad (\text{B.7})$$

Microsoft has \$2.10 of liquid assets for every dollar that it owes over the next year.

2. **Quick ratio.**

$$\text{Quick ratio}_t = \frac{\text{Cash}_t + \text{Short-term investments}_t + \text{Receivables}_t}{\text{Current liabilities}_t} = \frac{168,377}{88,657} = 1.9 \quad (\text{B.8})$$

(The “Cash” line on the balance sheet includes both cash and short-term investments.) Removing inventory, which is more difficult to convert to cash, than receivables or cash itself, reduces the ratio to 1.9. Microsoft has \$1.90 of (very) liquid assets for every dollar that it owes over the next year.




Figure B.6: Liquidity Measures

Figure B.6 presents the current and quick ratios for Microsoft and its competitors. All three companies have current ratios in excess of one for the entire period, suggesting that each has sufficient liquid assets to cover its current obligations. Of course, current assets are not the only means that firms have to pay their current obligations. They can also use earnings, as well as liquidating longer-lived assets, such as PP&E.

The differences in liquidity measures across firms again highlight different business models, as well as different liquidity management policies. Amazon runs a relatively tight ship compared to Microsoft and Google, keeping fewer liquid assets on its balance sheet compared to its current liabilities. It's also important to note the role of cash on these ratios. Each company has a significant amount of cash on its balance sheet, in excess of \$130 billion for Microsoft and Alphabet and almost \$100 billion for Amazon.

2.4.5 Working Capital Management

Working capital was defined earlier as current assets minus current liabilities.⁷ Closely related to liquidity, working capital is critical for companies to ensure they have enough cash to continue short-term operations and, by extension, long-term growth.

Unlike our liquidity measures, working capital ratios give us a sense of how the company manages its liquidity, as opposed to how much it has on hand. There are several key measures on which companies focus.

1. **Days sales outstanding (DSO)**. Also known as **days receivable**, **accounts receivable days** and **accounts receivable collection period**.

$$\begin{aligned}
 DSO_t &= \frac{(\text{Receivables}_{t-1} + \text{Receivables}_t)/2}{\text{Sales}_t} \times 365 & (\text{B.9}) \\
 &= \frac{(32,011 + 38,043)/2}{168,088} \times 365 \\
 &= 76.1
 \end{aligned}$$

It takes, on average, 76.1 days for Microsoft to collect money from its customers. This is a relatively long time, reflecting Microsoft's large business-to-business (B2B) component of sales. Most consumer transactions occur in less than 30 days.

A little intuition for what equation B.9 is doing. Imagine our customers buy goods from us on credit with 90 day terms; i.e., they pay us 90 days after buying the good. When we make our sales on the first day of the year, we won't collect the cash from these sales for another 90 days, on day 91 of the year. Same for sales made on day 2. We won't see the cash from those sales until day 92. And so on.

At the end of the year, we've sold product to our customers every day and collected cash from our customers for each day of sales *except* the last 90 days of the year. (Those sales won't be collected on until the first 90 days of next year.) So, at the end of the year we're still waiting to collect on 90 days of sales for the year out of the 365 days in the year. Mathematically,

$$\text{Accounts receivable}_t = \frac{DSO_t}{365} \times \text{Sales}_t$$

Solving this for DSO gets us the equation B.9 with the average receivables replaced by the year t receivables. (Recall, we average balance sheet items to recognize that the sales were generated by assets throughout the year.)

⁷Often we want to focus only on the current assets and liabilities involved in the operations of the company. **Operating working capital** excludes excess cash from current assets and any short-term debt or long-term debt due from current liabilities.

Some comments. First, the denominator should only contain sales that were made on credit, as opposed to cash. Unfortunately, companies don't break out sales into credit and cash components in their public filings. Unless you work at the company, and in a financial capacity, this distinction is unknowable.

Second, this computation assumes that sales are uniformly distributed throughout the year. In other words, the company sells the same amount each quarter. For companies with a highly seasonal business, equation B.9 can give a very misleading estimate of how long it takes to collect from customers. To illustrate the problem, consider a company that does all its business on Christmas of each year and has a December fiscal year end. Its end-of-year receivables will equal its sales for the year. Equation B.9 would imply that its DSO is equal to 365 when in fact it may require far fewer days to collect. Lesson: Be careful using equation B.9 with companies that have seasonal sales.

2. **Days payable outstanding (DPO).** Also known as **days payable**, **accounts payable days** and **days of accounts payable outstanding**.

$$\begin{aligned} DPO_t &= \frac{(\text{Payables}_{t-1} + \text{Payables}_t)/2}{\text{Cost of sales}_t} \times 365 && \text{(B.10)} \\ &= \frac{(12,530 + 15,163)/2}{40,546} \times 365 \\ &= 124.6 \end{aligned}$$

It takes, on average, 14.6 days for Microsoft to pay its suppliers.

The intuition for DPO is the same as the for DSO. As such, it is subject to the same limitations. Specifically, we should only be using credit purchases in the denominator, information that is rarely made available by the company. Likewise, days payable is not accurately represented by equation B.10 for highly seasonal business with non-uniform purchasing activity throughout the year.

3. **Days inventory outstanding (DIO).** Also known as **days inventory** and **days of sales held in inventory**.

$$\begin{aligned} DIO_t &= \frac{(\text{Inventory}_{t-1} + \text{Inventory}_t)/2}{\text{Cost of sales}_t} \times 365 && \text{(B.11)} \\ &= \frac{(1,895 + 2,636)/2}{40,546} \times 365 \\ &= 20.4 \end{aligned}$$

It takes, on average, 20.4 days for Microsoft to sell its inventory.

Again, DIO has a similar intuition and limitations as DSO and DPO.

4. Cash conversion cycle (CCC). Also known as **trade cash cycle**.

$$CCC_t = DSO_t + DIO_t - DPO_t = 76.1 + 20.4 - 124.6 = -28.2 \quad (\text{B.12})$$

The CCC measures the time from when the company must pay cash to when it collects cash. Put differently, it measures how long the company has to finance its supplier purchases. The larger this number, the longer the company has to finance (i.e., find money) its purchases.

Microsoft is in the relatively uncommon situation in its CCC is negative. Microsoft is using its suppliers to finance its operations. This is a great position to be in because the company is essentially getting interest-free loans from its suppliers. This situation is often limited to companies with strong bargaining positions over its customers or suppliers (e.g., Walmart) in which the company can push for shorter payment terms (low DSO) from its customers and longer payment terms to its suppliers (high DPO).

Of course, there is a tradeoff. Asking customers to pay too quickly may lead to customer loss, hurting sales. Likewise, asking suppliers to wait too long for payment may lead to suppliers exiting the market.

Figure B.7 presents the working capital metrics for Microsoft and its competitors. The differences across companies highlights the different business models.

The DSOs for Microsoft are all above 75 days. For Alphabet, DSOs are between 50 and 60 days. And, for Amazon, DSO are all under 25 days. Why the large differences? Different business models! Yes, Microsoft has significant retail arm in which customers are paying by credit card and even cash (in their retail stores before they were closed). But, most of their business is B2B, often involving longer-term contracts and payment terms. Amazon on the other hand has a massive online retail business whose revenue is generated almost entirely by credit card sales. While this would suggest an even lower DSO than 25 days, Amazon also has a non-trivial B2B business in cloud services and other areas.

Inventory is efficiently managed with DIO under 30 days. The particularly small DIO for Alphabet is largely an artifact of them having relatively little inventory - approximately \$1 billion on sales over \$250 billion. Microsoft's retail arms selling physical goods (e.g., Xbox, Surface laptops and accessories) is able to move product in less than a month. Amazon's figures, which are similar to Microsoft's, are perhaps most impressive despite them having the longest days in inventory. The scale and scope of their retail operations is incredibly efficient.




Figure B.7: Working Capital KPIs

There are also important distinctions among the companies in how they engage with their suppliers. Microsoft and Alphabet have similar terms, approximately four months. Amazon pays its suppliers in almost half that time. Of course, these differences also line up with differences in DSO. Amazon is much quicker to collect from their customers than both Microsoft and Alphabet.

In sum, working capital management is a critical aspect of corporate financial management. It keeps the lights on so to speak. It also provides additional channels for companies to create value that don't receive as much attention as sales growth.

2.4.6 Credit Risk

Credit risk refers to the risk that a company will not pay its obligations on time and in full. It affects firms with debt and therefore most firms.⁸ Some common credit risk KPIs include the following.

1. **Debt-to-ebitda ratio.** Also known as the **leverage ratio**.

$$\text{Debt-to-ebitda}_t = \frac{\text{Total debt}_t}{\text{EBITDA}_t} = \frac{67,775}{81,602} = 0.8x \quad (\text{B.13})$$

⁸As December 2021, 92% of all publicly traded firms had some debt on their balance sheet. An even larger fraction had some debt *and* leases. TODO: DOUBLE CHECK

Microsoft had a 2021 debt-to-ebitda ratio of 0.8x, read as debt equal to 0.8 times ebitda. Put differently, for every dollar of cash operating earnings, Microsoft has \$0.80 of debt meaning Microsoft could pay off all of its outstanding debt with one year's worth of earnings. As such, 0.8x is a relatively low value for the debt-to-ebitda ratio, indicative of a company with relatively little credit risk.

2. Interest coverage ratio.

$$\text{Interest coverage}_t = \frac{\text{EBITDA}_t}{\text{Interest expense}_t} = \frac{81,602}{-1,186} = -59.0\mathbf{x} \quad (\text{B.14})$$

The interest coverage ratio measures how many dollars of cash operating earnings a company has for each dollar of interest expense it faces. That Microsoft has a negative coverage ratio is an artifact of it having more interest income from its cash than interest expense from its debt. A somewhat more sensible figure can be found in 2016 when Microsoft's interest coverage ratio was 46.8x, meaning it had \$46.80 of earnings for each dollar of interest it owed. This coverage ratio is consistent with Microsoft's debt-to-ebitda ratio in that the company can easily manage its debt obligations with its operating earnings.

3. Debt service coverage ratio.

$$\begin{aligned} \text{Debt service coverage}_t &= \frac{\text{EBITDA}_t}{\text{Interest expense}_t + \text{Debt due this year}_t} \quad (\text{B.15}) \\ &= \frac{81,602}{-1,186 + 8,072} \\ &= 11.9\mathbf{x} \end{aligned}$$

The debt service coverage ratio expands on the interest coverage ratio to account for necessary principal payments. Microsoft has \$11.90 of cash operating earnings for each dollar of interest and principal it owes in 2021 - more than enough.

4. Debt-to-total capitalization ratio. Also known as the **leverage ratio**.

$$\begin{aligned} \text{Debt-to-total capitalization}_t &= \frac{\text{Total debt}_t}{\text{Total debt}_t + \text{Book equity}_t} \quad (\text{B.16}) \\ &= \frac{67,775}{67,775 + 141,988} \\ &= 0.323 \quad (\text{B.17}) \end{aligned}$$

The debt-to-total capitalization ratio measures the fraction of the firm's financing to come from debt as opposed to equity. Almost one third of Microsoft's financing has come from debt as opposed to common equity.

5. **Net debt-to-total capitalization ratio.**.. Also known as the **net leverage ratio**.

$$\begin{aligned} \text{Net debt-to-total capitalization}_t &= \frac{\text{Total debt}_t - \text{Cash}_t}{\text{Total debt}_t - \text{Cash}_t + \text{Book equity}_t} \quad (\text{B.18}) \\ &= \frac{67,775 - 130,334}{67,775 - 130,334 + 141,988} \\ &= 0.788 \quad (\text{B.19}) \end{aligned}$$

Cash in this formula refers to cash and short-term investments. The net debt-to-total capitalization ratio recognizes that firms can often use their cash holdings to pay down debt. Occasionally, you'll hear practitioners talk about cash as **negative debt** and debt minus cash as **net debt**. There are some important subtleties associated with treating cash as negative debt that are discussed in the financial policy chapter.

Microsoft is negatively levered. It has more cash than debt on its balance sheet. Thus, while the leverage ratio of 32.3% suggests a moderately levered firm facing some degree of credit risk, the large negative net leverage ratio shows the ideal borrower with virtually no credit risk. (The debt is backed by more than twice as much cash!)

Figure B.8 presents debt-to-ebitda, debt-to-total capitalization, and net debt-to-total capitalization measures for Microsoft and its competitors. The coverage ratios are less informative because all three companies have negative net interest expense; they're earning more in interest income than they are paying in interest expense.

All three companies face little credit risk. The figures illustrate why. Each has significant earnings relative to their debt, especially Amazon. Further, while the debt-to-capitalization ratios show Microsoft and Amazon as moderately levered, the *net* debt-to-capitalization ratios show these results are misleading. Each company is sitting on an enormous pile of cash, which is further reassuring for creditors.

We should also mention that liquidity and credit risk are closely related. Liquidity problems are the precursor to solvency problems (i.e., bankruptcy). Thus, it's not surprising to see bank loan contracts containing **covenants** (i.e. restrictions) requiring borrowers maintain different levels for their liquidity and credit risk metrics. For example, loan contracts might require a company maintain a current ratio above 0.5, a debt-to-ebitda ratio below 4.0, and an interest coverage ratio above 1.0.

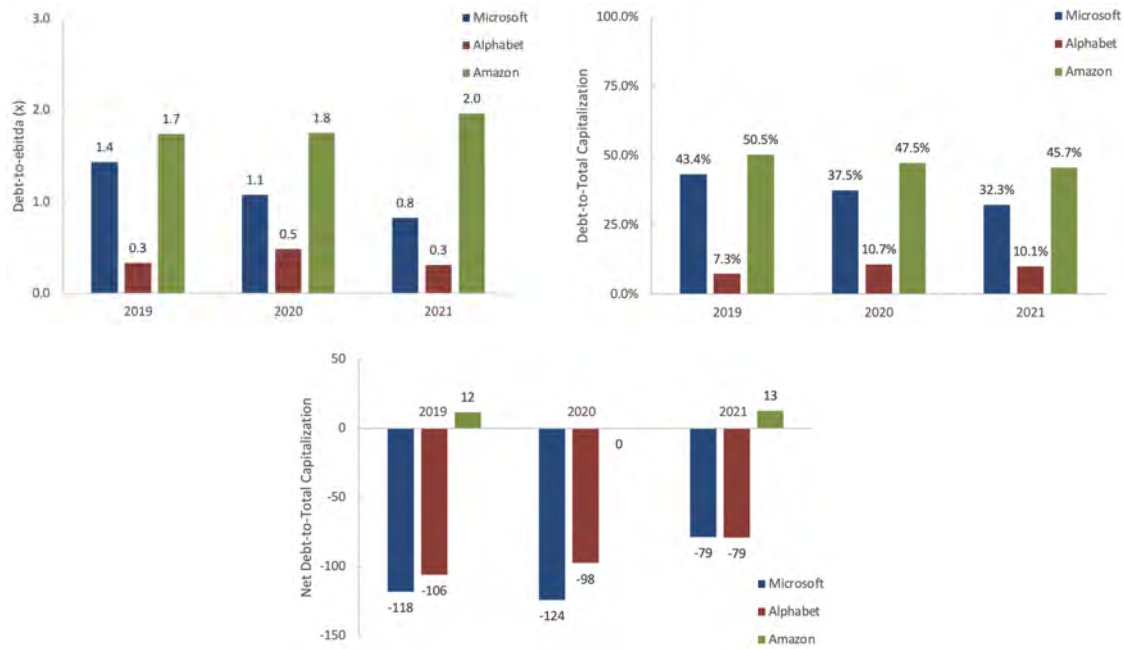


Figure B.8: Credit Risk KPIs

2.4.7 DuPont Analysis

The **DuPont analysis/formula/equation** gives us a way of understanding what's behind the return on equity (ROE) by decomposing ROE into its components.⁹

$$\underbrace{\frac{\text{Net income}}{\text{Equity}}}_{ROE} = \underbrace{\frac{\text{Net income}}{\text{Sales}}}_{\text{Net margin}} \times \overbrace{\frac{\text{Sales}}{\text{Assets}}}_{ROA}^{\text{Asset turnover/utilization}} \times \underbrace{\frac{\text{Assets}}{\text{Equity}}}_{\text{Financial leverage}} \quad (\text{B.20})$$

This decomposition shows that the way to drive shareholder returns, at least in book terms, is to (i) increase profitability (net margin), (ii) utilize assets more efficiently (increase turnover), or (iii) increase risk (increase financial leverage). Of course, these three channels are all closely related so changing one often affects another. When combined with historical data, DuPont analysis can provide a powerful tool for understanding and diagnosing the performance of a firm.

Figure B.9 presents the results of the Du Pont analysis applied to Microsoft. The left axis measures the net margins in percent, as represented by the blue bars. The right axis measures the financial leverage and asset turnover, as represented by the green and red bars, respectively. The annual ROEs are presented at the top of each group of bars.

⁹The name comes from the DuPont Corporation, the company at which it was invented.




Figure B.9: Du Pont Analysis of Microsoft's ROE

The decline in ROE between 2019 and 2020 was due to a slight decline in net margin and a more substantial decline in financial leverage. The increase in ROE from 2020 to 2021 was due to a large increase in net margin that more than offset the continuing decline in financial leverage. Over this period, Microsoft was reducing its reliance on debt financing - delevering - which reduces risk and therefore the return on equity. Counteracting this effect was an increase in the profitability of the company as seen by the changing net margin. Asset turnover was stable over this period suggesting that the efficiency with which Microsoft operated its asset was unchanged; each dollar of assets generated \$0.50 of sales.

Of course, we can dig even deeper. For example, why exactly did profitability increase so much from 2020 to 2021? Table 8 shows the annual growth rates for each line item on its P&L.

We see that sales grew by 17.5% but cost of sales grew by even more (21.7%), implying a *lower* gross margin. This rules out sales growth as a driver of the increased profitability. SG&A, on the other hand, grew at a very modest rate of 4.5% and depreciation and amortization contracted by 8.5%. We also see a very large increase in other income, which also contributed to an increased bottom line. When we compare the scale of these three line items, SG&A is by far the largest. Consequently, its relatively modest growth compared to sales growth is what is largely responsible for the increased net margin in 2021 with

	2021
Sales	17.5%
Cost of sales	21.7%
Gross profit	16.3%
SG&A	4.5%
EBITDA	24.2%
Depreciation & amortization	(8.5%)
EBIT	32.0%
Other expenses (income)	1,440.3%
Pre-tax income	34.1%
Taxes	12.3%
Net income	38.4%

Table 8: Microsoft 1-Year Growth Rates from 2020 to 2021

depreciation and amortization and other income playing secondary roles.

Of course, we can continue this process even further by asking: Why did SG&A grow so slowly compared to sales? Why did depreciation and amortization expense shrink? And, why is Microsoft earning so much more in interest income? A careful reading of their 10-k can provide some clues. More importantly, this exercise illustrates what we stated at the beginning of this chapter. Financial statement analysis is more likely to help us ask the right questions of managers as opposed to providing us with the final answer.

2.4.8 Rule of 40

The “rule of 40” is an ad hoc rule used to identify firms on a positive growth trajectory. Investors in the growth equity (i.e., private equity focused on smaller, growth oriented firms) space are particularly fond of this rule which seeks to identify firms whose sales growth and ebitda margin sum to more than 40%. Of course, there are a variety of ways for firms to achieve a sum of 40%.

Firms can have balanced growth and profitability. Alternatively, firms can be growing rapidly but not quite yet profitable. Finally, firms can be slowing in their growth but extremely profitable. The rule of 40 attempts to recognize this tradeoff between sales growth and profitability. Why 40? Not clear, but the tradeoff does suggest an empirical study in the future if one doesn’t already exist.


	2019	2020	2021
Microsoft			
Sales growth (%)	14.0	13.6	17.5
EBITDA margin (%)	43.4	46.0	48.5
Sum (%)	57.4	59.6	66.1
EV-to-EBITDA (x)	17.7	22.4	24.2
Alphabet			
Sales growth (%)	18.3	12.8	41.2
EBITDA margin (%)	29.5	30.1	35.4
Sum (%)	47.8	42.9	76.5
EV-to-EBITDA (x)	7.4	8.5	8.8
Amazon			
Sales growth (%)	20.5	37.6	21.7
EBITDA margin (%)	13.0	12.5	12.6
Sum (%)	33.4	50.1	34.3
EV-to-EBITDA (x)	25.6	34.0	29.0

Table 9: Rule of 40 Calculations and Market Multiples

Table 9 presents the data for the rule of 40 analysis. The sum row is the sum of the sales growth and EBITDA margin. The EV-to-EBITDA row contains the enterprise value-to-EBITDA ratio. **Enterprise value** equals the debt less cash plus market capitalization. It is a measure of the ongoing value of the company were it to distribute all of its cash to investors. The enterprise value-to-EBITDA ratio reveals how much investors - debt and equity - are willing to pay for each dollar of operating income.

We can see that all but two of the Sum values for Amazon are greater than 40. A closer look reveals that for most firm-years, most of the action is in profitability. Though some years we see impressive growth, especially for Amazon in all three years and Alphabet in 2021.

Figure B.10 plots the data. The data labels present the enterprise value-to-EBITDA ratio. The dashed line corresponds to the rule of 40 boundary. Data points below and to the left of the dashed line have sales growth and EBITDA margin sums less than 40%, those above and to the right have sum greater than 40%. (Very) loosely speaking, the enterprise value-to-EBITDA ratio for each firm tends to increase as its data points move Northeast, suggesting that the investors place greater value on firms with more combined sales growth and profitability.



*Data labels = Enterprise value-to-EBITDA

Figure B.10: Rule of 40 Analysis

B.5 Key Ideas

- Financial statements are how financial information is organized and presented. They are “backward looking” in that they are based on historical transactions values, as opposed to “forward looking” market values based on future cash flows and discount rates.
- The income statement (a.k.a., P&L, statement of operations) details the sales, expenses, and earnings over a time period. Accrual accounting means that sales and expenses are recognized when a transaction occurs, as opposed to when money is received or paid.
- The balance sheet presents a snapshot at a point in time of what the companies owns (assets) and what the company owes or how it purchased those assets (liabilities and shareholders equity).
- The statement of cash flows details the actual money flowing in and out of a company over a time period.
- Financial statement analysis provides insights into a company by revealing which questions to ask.

- Key performance indicators (KPIs) need context to be interpreted. This context can come from historical data, competitor's data, or expectations.

Appendix C

Spreadsheet Tips

This appendix presents a collection of useful spreadsheet functions. It is *not* meant to be exhaustive. In fact, we only scratch the surface here of what Excel is capable. For more details, there are countless books exploring the capabilities of Excel. And, though we'll illustrate concepts using Excel, most everything we mention here translates to Google's Sheets or other spreadsheet products.

C.1 Present Values - Constant Cash Flows

Figure C.1 is a screen shot of an annotated, Excel spreadsheet that computes the present value of the cost of college. Blue font corresponds to hard-coded numbers, i.e., numbers typed into the cell. Black font corresponds to numbers computed from a formula. Rows 4 and 6 correspond to the timeline of cash flows as originally displayed in figure 2.1. Cell G2 contains our 5% opportunity cost.

We've computed the present value of the cash flows using four different approaches, all of which give us the same answer, \$297,859.84. Doing so allows us to discuss some of Excel's capabilities. Which approach we use depends upon the circumstances.

The first approach discounts each cash flow to find their present value and then sums. The present value of each cash flow is in row 9 with the corresponding Excel formula just above each cell. The second approach recognizes the cash flows as an annuity and applies our annuity formula (equation 2.7). Because the first cash flow occurs today - period 0 - the annuity formula is applied to the three future cash flows in years one through three.

The third approach exploits Excel's **NPV function**, which takes as its arguments a discount rate and a sequence of cash flows (one or more). Unfortunately, the NPV function

							$\frac{80,000}{(1+0.05)^0}$	$\frac{80,000}{(1+0.05)^1}$	$\frac{80,000}{(1+0.05)^2}$	$\frac{80,000}{(1+0.05)^3}$		
	A	B	C	D	E	F	G	H	I	J	K	L
1												
2		Opportunity cost										
3												
4		(\$000s)						0	1	2	3	4
5												
6		Cash flows						80,000.0	80,000.0	80,000.0	80,000.0	
7												
8		Cost of college						$=H6/(1+\$G\$2)^H4$	$=I6/(1+\$G\$2)^I4$	$=J6/(1+\$G\$2)^J4$	$=K6/(1+\$G\$2)^K4$	
9		Present values at t=0						80,000.00	76,190.48	72,562.36	69,107.01	
10		Sum of present values						297,859.84	$=SUM(H9:L9)$			
11		Annuity formula						297,859.84	$=H6+(I6/G2)*(1-(1+G2)^(-K4))$			$80,000 + \frac{80,000}{0.05} \times (1 + (1 + 0.05)^{-3})$
12		Excel NPV function						297,859.84	$=H6+NPV(G2,I6:K6)$			
13		Excel PV function						297,859.84	$=-PV(G2,L4,H6,-1)$			

Figure C.1: Computing Present Values in Excel

does not compute the NPV as we know it in finance because it discounts *all* of the cash flows that are passed to it. As such, we have to break out today's payment in cell H6, like we did with our annuity formula above, and apply the NPV function only to the future cash flows in periods one through three. What's nice about this NPV function is that the cash flows don't have to be the same. They do; however, have to be equally spaced in time and subject to the same discount rate, r .

The final approach uses Excel's **PV function**, which computes the present value of an annuity. It's less flexible than the NPV function because the cash flows must all be the same. However, it can handle cash flows that start today. We haven't broken out today's payment in cell H6. We simply passed to the function the discount rate (G2), the number of payments (L4), the payment (H6), and the value "1" to indicate that the first cash flow comes today. If the first cash flow came one period in the future, we would pass the value "0." We've put a negative sign in front of the function because Excel returns the negative of the present value as it assumes a sequence of payments.¹

C.2 Present Values - Growing Cash Flows

Figure C.2 illustrates present value computations for cash flows growing at a constant 3% rate. Discounting each cash flow and summing, and use of the NPV function are unchanged from the constant cash flow example. Our annuity formula is modified for growing cash flows (equation 2.7). To use the Excel PV function, we replace the nominal growth rate with the real growth rate (equation 2.5). Again, all four approaches generate the same result.

¹One argument of the PV function is left blank. See the Excel documentation for details.

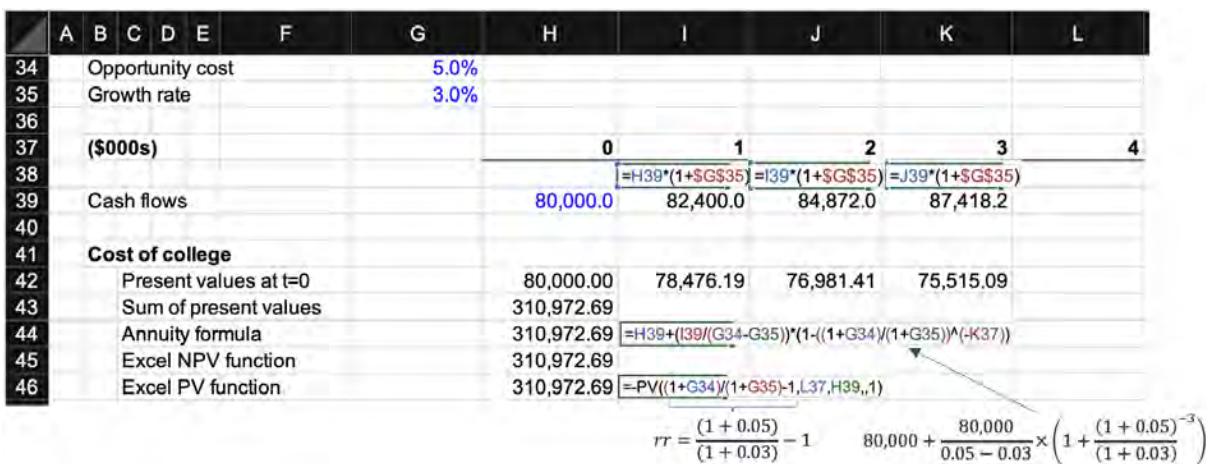


Figure C.2: Computing Present Values in Excel

A word of caution: Excel’s NPV formula treats blank cells differently from cells with a zero. Consider figure C.3, which shows two different NPVs for what appear to be the same set of cash flows. The first assumes the blank cell doesn’t exist and incorrectly discounts the cash flows. The lesson is to avoid blank cells in our timelines and instead use zeros.




Figure C.3: Excel’s NPV Function Treatment of Blank Cells

C.3 Annuity Cash Flows

Figure C.4 presents annuity cash flow calculations for Sophie’s annual savings requirement assuming the present value of her nest egg is \$807,232.74. Cell I46 uses the annuity cash flow formula from equation 2.2 to show that Sophie must save \$34,922.81 every year, for 40 years beginning next year if she is to meet her savings goal at the start of retirement. Cell I47 uses Excel’s **PMT function** to arrive at the same result. The PMT function takes

as arguments the discount rate (I39), the number of periods (hard coded as 40), and the present value (I45). The negative sign in front of the PMT function converts the result to a positive number.

If Sophie wants to save 2% more each year, we can use the annuity cash flow for a growing annuity (equation 2.9) as in cell I48. We get the same result in cell I49 using the PMT function with two modifications. First, we use the real rate of return $(1+0.03)/(1+0.02) - 1 = 0.98\%$. This will return the savings amount as of today - period 0. To get the amount for next year, we have to multiply by one plus the growth rate.

	A	B	C	D	E	F	G	H	I
39			Discount rate						3.0%
40			Growth rate						2.0%
41									
42			Age						30
43			Period						0
44									
45			Present value nest egg						807,232.74
46			Constant annual savings (annuity cash flow formula)						34,922.81 =I45*I39/(1-(1+I39)^-40)
47			Constant annual savings (Excel PMT function)						34,922.81 =-PMT(I39,40,I45,,0)
48			Growing annual savings (annuity cash flow formula)						24,983.19 =I45*(I39-I40)/(1-((1+I39)/(1+I40))^-40)
49			Growing annual savings (Excel PMT function)						24,983.19 =-PMT((1+I39)/(1+I40)-1,40,I45,,0)*(1+I40)

Figure C.4: Computing Annuity Cash Flows

C.4 Future Values

Figure C.5 is a screen shot of an annotated, Excel spreadsheet that computes the *future* value of the cost of college as of four years from today. Rows 4 and 6 correspond to the timeline of cash flows as originally displayed in figure 2.3. Cell G2 contains our 5% opportunity cost.

We've computed the future value of the cash flows using three different approaches, all of which give us the same answer, \$362,050.50. The first approach is a brute force approach that separately computes the future value of each cash flow and then sums. The second approach takes the present value computed above, \$297,859.84, and computes its future value. (Rows 8 through 15 are hidden to manage the size of the screen shot. The present values of the cash flows are contained in cells H10 through H13.)

The final approach uses Excel's **FV function**, which computes the future value of an annuity. We pass to the function the discount rate (G2), the number of payments (L4), the payment (H6), and the value "1" to indicate that the first cash flow comes today. If the

	A	B	C	D	E	F	G	H	I	J	K	L
1												
2							5.0%					
3												
4								0	1	2	3	4
5												
6								80,000.0	80,000.0	80,000.0	80,000.0	
7												
16								$=H6*(1+G$2)^{(L$4-H4)}$	$=I6*(1+G$2)^{(L$4-I4)}$	$=J6*(1+G$2)^{(L$4-J4)}$	$=K6*(1+G$2)^{(L$4-K4)}$	
17								97,240.50	92,610.00	88,200.00	84,000.00	
18												$=SUM(H17:K17)$
19												$=H11*(1+G2)^{L4}$
20												$=-FV(G2,L4,H6,,1)$

Figure C.5: Computing Future Values in Excel

first cash flow came one period in the future, we would pass the value “0.” We’ve put a negative sign in front of the function because Excel returns the negative of the future value as it assumes a sequence of payments.²

C.5 Estimating an Internal Rate of Return

Figure C.6 presents the Excel worksheet where we derived the auto lease interest rate. The timeline in rows 35 and 37 show the months and cash flows, respectively. We’ve hidden months four through 35 (columns N through AS) to ease the presentation. The lease is just like a loan in that we are receiving money exchange for future repayment.

	A	B	C	D	E	F	G	H	I	J	K	L	M	AT
35										0	1	2	3	36
36														
37										\$109,000.0	(\$4,820.0)	(\$4,820.0)	(\$4,820.0)	(\$4,820.0)
38								2.77%	$=IRR(J37:AT37)$					
39								33.19%	$=H38*12$					
40								38.74%	$=(1+H39/12)^{12}-1$					

Figure C.6: Computing Loan Interest Rates in Excel Using IRR

To find the periodic interest rate of the loan, we use the **IRR function** which computes the **internal rate of return (IRR)**. The IRR function finds the one discount rate such that the net present value of all the cash flows equals zero. Mathematically, the IRR is found by solving the following equation, which requires the use of a computer for any setting with

²One argument of the FV function is left blank. See the Excel documentation for details.

more than two or three periods.

$$NPV = 0 = CashFlow_0 + \frac{CashFlow_1}{(1 + IRR)} + \frac{CashFlow_2}{(1 + IRR)^2} + \dots + \frac{CashFlow_T}{(1 + IRR)^T}$$

The only argument that needs to be passed to the IRR function is the cash flows. You can also pass a guess at the IRR, but this is rarely needed. To get the APR, we multiply the periodic rate by the compounding frequency, 12 in this instance. The loan APR is computed using equation (3.2). To use the IRR function, the cash flows must be ordered and equally-spaced in time.

Instead of using the IRR function, we could use the **Goal Seek** feature which can be found under the “Data-What if Analysis” menu. This function will set the value of an equation to any number we want by changing one of the inputs to that equation. For example, if we discount our lease payments by 3.0% - chose arbitrarily - the lease NPV equals

$$109,000 + \frac{-4,820}{(1 + 0.03)} + \frac{-4,820}{(1 + 0.03)^2} + \dots + \frac{-4,820}{(1 + 0.03)^{36}} = \$3,768.54.$$

We can use Goal Seek to set the NPV equal to zero by changing the periodic interest rate, as illustrated in Figure C.7. After clicking OK, Excel will find the periodic interest rate that sets the NPV of the lease (cell H42) equal to zero by changing the number in cell H41.

	A	B	C	D	E	F	G	H	I	J	K	L	M	AT
34														
35														
36														
37														
38														
39														
40														
41														
42														

Figure C.7: Computing Loan Interest Rates in Excel Using Goal Seek

C.6 Estimating a Modified Internal Rate of Return

As discussed in Chapter 6, sometimes there are multiple IRRs or non-real (complex) IRRs because of the nature of the cash flows. When this occurs, we can use the modified internal rate of return or **MIRR** Excel function to get a number, though the interpretation of the

Bids	0	1	2	3	NPV	IRR	MIRR
Juniper 2	50	-60	75		\$56.22	#NUM!	=MIRR(D39:G39,0.12,0.12)

Figure C.8: Computing Loan Interest Rates in Excel Using Goal Seek

number it produces is unclear. MIRR stands for **Modified Internal Rate of Return**. Its use is illustrated in Figure C.8.

The function has three arguments the first of which are the cash flows, which must be ordered from earliest to latest and equally-spaced in time. The second argument is the **finance rate**, which corresponds to the interest rate we pay on the cash outflows. This sounds like the cost of capital. The third argument is the **reinvest rate**, which is the expected return we earn on any cash inflows. Where the IRR assumes that the finance and reinvest rates are the same, the MIRR allows them to be different. From this information, MIRR produces its modified internal rate of return. See the technical appendix of chapter 6 for details on the computation that MIRR performs.

C.7 Working with Mortgages and the Goal Seek Function

Figure C.9 illustrates some common computations in a mortgage or, more generally, an amortizing loan setting (e.g., auto loan, student loan). The monthly mortgage payment is an annuity cash flow that can be computed using equation 2.2 or the **PMT function**, as discussed in the technical appendix of chapter 2. The PMT function takes as arguments the periodic interest rate (H11), the number of periods (H12), and the loan amount. The negative sign in front of the PMT function converts the result to a positive number.

The outstanding principal at any time during the life of the loan can be computed by discounting and summing all remaining cash flows using the annuity formula (equation 2.1) or the PV function, which we saw in chapter 2.

Rows 20 through 23 decompose the 60th mortgage payment into its interest and principal component. Cell H20 computes the interest component by multiplying the monthly interest rate by the outstanding principal as of the end of the 59th month (just after the 59th payment), latter of which is computed using the present value of an annuity equation (2.1) from chapter 2. Alternatively, we can use Excel's **IPMT function** to get the same result by passing the periodic interest rate (H11), the payment period (H17), the number of periods in

	A	B	C	D	E	F	G	H
8	*	Existing mortgage						
9		Mortgage principal						500,000.0
10		Mortgage APR						3.0%
11		Periodic interest rate						0.2500% <small>=H10/12</small>
12		Mortgage term (months)						360.0
13		Monthly mortgage payment (annuity equation)						\$2,108.02 <small>=(H9*H11)/(1-(1+H11)^(-(H12)))</small>
14		Monthly mortgage payment (PMT function)						\$2,108.02 <small>=PMT(H11,H12,H9)</small>
15								
16	*	Refi mortgage						
17		Refi period (month)						60.0
18		Outstanding principal (annuity equation)						\$444,531.82 <small>=H13/H11*(1-(1+H11)^(-(H12-H17)))</small>
19		Outstanding principal (PV function)						\$444,531.82 <small>=PV(H11,(H12-H17),H13)</small>
20		Interest component of 60th payment						\$1,113.82 <small>=H11*(H14/H11)*(1-(1+H11)^(-(H12-(H17-1))))</small>
21		Interest component of 60th payment						\$1,113.82 <small>=IPMT(H11,H17,H12,H9,0)</small>
22		Principal component of 60th payment						\$994.21 <small>=H14-H20</small>
23		Principal component of 60th payment						\$994.21 <small>=PPMT(H11,H17,H12,H9,0)</small>

Figure C.9: Computing Mortgage Quantities in Excel

the mortgage (H12), and the original mortgage principal (H9). The last optional argument indicates that the first payment occurs one period from today.

We can get the principal component of the 60th mortgage payment by simply subtracting the interest component from the mortgage payment - cell H22. Or, we can use Excel's **PPMT function** which requires the periodic interest rate (H11), the payment period (H17), the number of periods in the mortgage (H12), and the original mortgage principal (H9).

C.8 Basic Statistical Functions

Figure C.10 is a screen shot of an annotated, Excel spreadsheet that presents four monthly returns to Exelon Corp. Blue font corresponds to hard-coded numbers, i.e., numbers typed into the cell. Black font corresponds to numbers computed from a formula.

The figure shows how to manually compute the average, variance, and standard deviation using equations 8.2 and 8.3, and the corresponding built-in Excel functions **AVERAGE**, **VAR.S**, and **STDEV.S**.³ Also shown is the Excel function **COUNT**, which counts the number of cells in a highlighted range containing numbers, and the **SQRT** function, which computes the square root of a positive number.

³The “.S” in the variance and standard deviation functions identify the calculation of the **sample** variance and standard deviation, as opposed to the **population** variance and standard deviation. The latter two quantities can be computed using the **VAR.P** and **STDEV.P** Excel functions. The sample versions divide the sum of squared deviations by the number of observations minus one, the population versions divide by the number of observations.

	A	B	C	D
1				
2			Exelon	Deviation
		Date	Returns (%)	Squared
3		10/31/2000	-0.46	0.87
4		11/30/2000	10.19	94.38
5		12/29/2000	5.98	30.31
6		1/31/2001	-13.82	204.14
7		Sum	=SUM(C3:C6) 1.89	329.70 =SUM(D3:D6)
8		Average	=C7/COUNT(C3:C6) 0.47	
9		Average	=AVERAGE(C3:C6) 0.47	
10		Variance	=VAR.S(C3:C6) 109.90	109.90 =D7/(COUNT(D3:D6)-1)
11		Standard Deviation (Volatility)	10.48	10.48 =SQRT(D10)

Figure C.10: Some Excel Statistical Functions

Figure ?? shows the use of correlation (**CORREL**) and covariance (**COVARIANCE.S**) functions in Excel. Each function takes two arguments corresponding to the series for which we want to find the correlation or covariance. The example in the figure explores the correlation and covariance between Microsoft's (MSFT) monthly stock returns found in cells C5 through C65 and Walmart's monthly stock returns found in cells D5 through D65.

	G	H	I	J
3		Monthly		
4		MSFT	WMT	Risk-Free
5	Mean	1.89%	0.59%	0.11%
6	SD	6.45%	4.63%	0.00%
7	Corr	(0.07)	=CORREL(C5:C65,D5:D65)	
8	Cov	(0.00)	=COVARIANCE.S(C5:C65,D5:D65)	
9	Sharpe	0.28	0.10	

Figure C.11: Correlation and Covariance Functions

C.9 Dealing with Circular References

In chapter 11, we encountered situations in which two quantities depended on one another. For example, in the project financing example, the enterprise value was a function of the

weighted average cost of capital (WACC) and the WACC was a function of the enterprise value. In a spreadsheet, this situation leads to a **circular reference**. Figure C.12 illustrates the situation.

	B	C
124	Tax rate	21.00%
125	NOPAT	10.00
126	Cash	20.00
127	Debt	50.00
128	Debt cost of capital	5.00%
129	Operating cost of capital	10.00%
130	WACC	9.69% <small>=C129-(C127-C126)/C131*C124*C128</small>
131	Enterprise value	103.15 <small>=C125/C130</small>

Figure C.12: Circular Reference

Cell C130 contains the WACC equation,

$$r^{WACC} = r^O - \frac{D - C}{\text{Enterprise value}} \times \text{Tax Rate} \times r^D, \quad (\text{C.1})$$

where r^O is the operating cost of capital (C129), D is debt (C127), C is cash (C126), and r^D (C128) is the debt cost of capital. The WACC is also a function of enterprise value (C131). Cell C131 contains the enterprise value equation, which we assume can be evaluated as a perpetuity.

$$\text{Enterprise value} = \frac{NOPAT}{r^{WACC}} \quad (\text{C.2})$$

The NOPAT value is contained in cell C125 and the WACC in C130. The circular reference refers to cells C130 and 131; they refer to one another. To get Excel or Google Sheets to evaluate these two cells properly, we need to take the following steps (in order).

1. Turn on “Use iterative calculation.” In Excel for PCs, this feature can be found in Files-Options-Formula. In Excel for Mac, this feature can be found in Excel-Preferences-Calculation. In Google Sheets, this feature can be found in File-Spreadsheet settings-Calculation.
2. Enter an arbitrary number for one of the cells containing the circular reference. For example, we could enter 10% for the WACC (C130), or 100 for the enterprise value (C131).

3. Enter the correct formula for the other cell containing the circular reference. When you do this, make sure the result is a valid number.
4. Go back to step 2. and enter the correct formula. For example, if we entered 10% for the WACC to get the enterprise value, go back and enter the correct WACC formula, equation [C.1](#). If instead we had entered 100 for the enterprise value, go back and enter the correct enterprise value formula, equation [C.2](#).

After these steps, the spreadsheet will iteratively calculate both cells until the correct answers are found. If the program iterates without finding a solution, then there is most likely an error in one of the two equations.

Index

β , *see* beta

401-k, 608

403-b, 608

ABS, *see* asset-backed securitization

absolute priority rule, 234

account balance, 229

accounts payable, 631

 calculation, 155

 days, 148

accounts payable days, 646

accounts receivable, 630

 calculation, 154

 sales seasonality, 155

 days, 148

accounts receivable collection period, 645

accounts receivable days, 645

accrual accounting, 623

accure, 82

acquirer, 247

acquisition, 247

activist investors, 343

adjustable rate mortgage, 81

adjustable-rate preferred stock, 326

adjusted present value, 551

after hours trading, 332

after-tax expected return, 36

all-or-none commitment, 329

allocated costs, 150

amortization, 18

 in free cash flow, 142

 table, 18, 82

amortizing bonds, 233

amortizing loan, 77

annual percentage rate, 69

annual percentage yield, 73

annuity, 23

 annuity equivalent cash flow, 203

 cash flow, 25

 growing cash flow, 38, 42

 length of time, T , 28

 present value, 23

APR, *see* annual percentage rate, *see*
 absolute priority rule 234

APV, 551

APY, *see* annual percentage yield

arbitrage, 106, 275

 principle of no arbitrage, 299

arbitrageurs, 277

ARM, *see* adjustable rate mortgage

ask price, 242, 329

asset

 liquidation value, 148

 salvage value, 147

 usable life, 147

- asset turnover, 642
- asset utilization, 642
- asset-backed securitization, 292
- assets, 629
 - liquid, 630
 - marketable securities, 630
 - short-term investments, 630
- average tax rate, 142
- balance sheet identity, 633
- bank discount basis, 243
- bankruptcy costs
 - direct, 528
 - indirect, 528
- basis point, 90
- bequests, 26
- best efforts commitment, 329
- beta, 461
 - market, 465, 466
 - risk-free, 468
 - rolling, 471
- bid price, 242, 329
- bid-ask spread, 242
- Bloomberg adjustment, 490
- board of governors, 297
- bond
 - arbitrage, 275
 - order type, 330
 - yield vs. expected return, 288
- bond equivalent yield, 243
- bonds, 227, 232
 - amortizing, 233
 - bank discount basis, 243
 - bullet, 233
 - collateral, 234
 - convexity, 280
 - coupon, 233
 - taxes, 304
 - coupon rate, 261
 - credit ratings, 289
 - credit spread, 289
 - dealers, 242
 - dedication, 249
 - default recovery rate, 288
 - dollar value of a basis point (DV01), 281
 - duration
 - Macaulay, 283
 - Modified, 284
 - immunization, 286
 - inflation risk, 292
 - marketable, 232
 - maturity, 233
 - nominal, 293
 - original-issue discount (OID), 303
 - portfolio duration, 286
 - priced at discount, 268
 - priced at par, 268
 - priced at premium, 268
 - priority, 234
 - rating agencies, 290
 - reinvestment risk, 272
 - required yield, 243
 - secured, 234
 - seniority, 234
 - stripping coupons, 238
 - taxes on zero coupon, 303
 - term, 233
 - tradeable, 232
 - Treasurys, 236
 - yield, 243
 - yield-to-maturity, 243
 - YTM, 243

- zero-coupon
 - expected return, 258
- book values, 499
- bootstrapping, 266
- bullet bonds, 233
- C-corp, *see* C-corporation
- C-corporation, 575
- C-STRIPS, 238
- CAC, *see* customer acquisition costs
- CAL, *see* capital allocation line
- callable bonds, 239
- capital, 124
- capital allocation line, 421
- Capital Asset Pricing Model, 457
- Capital Asset Pricing model, 349
- capital budgeting, 127
- capital expenditures
 - in free cash flow, 144
- capital gain, 153
- capital gains
 - tax, 153
- capital loss, 153
- capital market line, 459, 479
- capital structure
 - predetermined, 550
 - target, 548
- CAPM, *see* Capital Asset Pricing Model,
see Capital Asset Pricing Model
- carrying value, 631
- cash accounting, 623
- cash conversion cycle, 647
- cash flow, 2
 - incremental, 47
 - inflow, 2
 - nominal, 31
 - outflow, 2
 - real, 31
- cash flow matching, 249
- cash flow statement, 634
- cash multiple, 138
- cash yield, 138
- cash-on-cash return, 138
- CCC, *see* cash conversion cycle
- CD, *see* certificate of deposit, *see*
 - certificate of deposit
- certificate of deposit, 227, 230
 - automatically renewable, 232
 - ladder, 232
- change in net working capital, *see* net
 - working capital investment
- change in working capital, *see* net
 - working capital investment
- churn rate, 208
- CLO, *see* collateralized loan obligation
- CLV, *see* customer lifetime value
- CML, *see* capital market line, *see* capital
 - market line
- COGS, *see* cost of goods sold
- collateral, 234, 277, 529
- collateralized loan obligation, 292
- collective action problems, 236
- comparable companies, 591
- compound interest, 70
- compounding, 9
 - frequency, 72
- comps, *see* comparable companies
- conflict of interest, 323
- consumer price index, 293
- continuation value, 363, 577
- continuous compounding
 - interest rate, 109
- convertible preferred equity, 326

- COR, *see* cost of revenue
- corporate governance, 323
- Corporate Social Responsibility, 128
- corporation, 575
- correlation, 403
- COS, *see* cost of sales
- cost of capital, 6, 84
 - unlevered, 511
- cost of equity capital, 126
- cost of goods sold, 624
- cost of revenue, 624
- cost of sales, 624
- counterparty risk, 230
- coupon bonds, 233
- coupon rate, *see* bonds
- covenant, 650
- covenants, 235, 532
- CPI, *see* consumer price index, *see*
 - consumer price index, 293
 - reference, 294
- credit card, 104
- credit ratings, 289, 290
- credit risk, 287
 - corporate, 648
- credit spread, 289
- creditors, 125
- CSR, *see* Corporate Social Responsibility
- consumer price index, 32
- current liability, 631
- current ratio, 644
- customer acquisition costs, 208
 - payback period, 208
- customer lifetime value, 207
- D&A, *see* depreciation and amortization
- days inventory, 646
- days inventory outstanding, 646
- days of accounts payable outstanding, 646
- days of sales held in inventory, 646
- days payable, 646
- days payable outstanding, 646
- days receivable, 645
- days sales outstanding, 645
- DCF, 139
- DDM, *see* dividend discount model
- debt
 - corporate, 632
 - long-term, 632
 - long-term due, 632
 - short-term, 632
- debt cost of capital, 127
- debt service coverage ratio, 649
- debt-to-capitalization ratio, 649
- debtholders, 125
- dedication, 249
- deductions, 605
- default, 105, 287
 - payment, 287
 - technical, 287
- default risk, 287
- defined benefit plan, 612
- defined contribution plan, 612
- deflate, 32
- deflation, 30
- demand deposit, 228
- deposit, 18
- depreciation
 - accumulated, 631
 - calculation, 151
 - in free cash flow, 142
 - straight-line, 626
- depreciation and amortization, 626
 - in free cash flow, 142

- tax shield, [144](#)
- dilution, [536](#)
- DIO, *see* days inventory outstanding
- discount bonds, [233](#)
- discount factor, [10](#)
- discount rate, [6](#)
- discounted cash flow, [139](#)
 - CLV relation to, [209](#)
- discounted payback period, [137](#)
- discounting, [8](#)
 - exchange rate for time, [8](#)
 - money's time unit, [8](#)
- Diversifiable risk, [407](#)
- diversification, [405](#)
 - dollar risk, [431](#)
- dividend discount model, [344](#)
- dividend payout ratio, [348](#)
- dividend yield, [333](#)
- dividends, [323](#), [335](#)
 - declaration date, [335](#)
 - ex-date, [335](#)
 - ex-dividend date, [335](#)
 - payment date, [335](#)
 - record date, [335](#)
- dollar risk, [431](#)
- dollar value of a basis point
 - bonds, [281](#)
- double taxation, [575](#)
- DPO, *see* days payable outstanding
- DSO, *see* days sales outstanding
- dual mandate of the fed, [297](#)
- dual-class shares, [325](#)
- DuPont
 - analysis, [651](#)
 - formula, [651](#)
- DV01, *see* bonds
- EAR, *see* effective annual rate
- earnings, [628](#)
 - operating, [628](#)
- earnings before interest after taxes, [141](#)
- earnings before interest and taxes, [142](#)
- earnings before interest taxes depreciation and amortization, [142](#)
 - in free cash flow, [142](#)
- earnings per share, [359](#)
- EBIAT, *see* earnings before interest after taxes
- EBIT, *see* earnings before interest and taxes, [628](#)
- EBITDA, *see* earnings before interest taxes depreciation and amortization, [628](#)
- effective annual rate, [70](#)
- effective tax rate, [142](#)
- elasticity, [161](#)
 - semi-elasticity, [284](#)
- embedded options, [239](#)
- enterprise value, [500](#), [654](#)
- Environmental, Social, and Governance, [128](#)
- EPS, *see* earnings per share
- equity, [322](#)
- equity cost of capital, [342](#)
- equityholders, [124](#)
- ESG, *see* Environmental, Social, and Governance
- Excel
 - FV, [660](#)
 - Goal Seek, [662](#)
 - IPMT, [663](#)
 - IRR, [661](#)
 - NPV, [657](#)

- PMT, 659, 663
- PPMT, 664
- PV, 658
- excess assets, 500, 630
- excess cash, 146
- excess interest, 526
- excess return, 482
- expectation, *see* statistical expectation
- expectations hypothesis
 - interest rate, 296
- expected debt return, 127
- expected return, 6
- expected stock return, 126
- expense
 - non-operating, 627
 - operating, 624
- expense ratios, 640
- expenses
 - in free cash flow, 142
 - overhead, 147
- face
 - bonds, 233
- factor models, 457
- FASB, 555
- FCF, *see* free cash flow
- FCFE, *see* free cash flow to equity
- Fed, *see* federal reserve
- federal reserve, 297
 - board of governors, 297
 - open market committee, 297
 - reserve banks, 297
- Financial Accounting Standards Board, 555
- financial constraint, 196
- financial constraints, 191
- financial distress, 528
- financial policy, 142
- financial statement analysis, 637
- firm commitment, 329
- firm value, 517
- firm-specific risk, 408
- Fitch, 290
- fixed cost, 625
- fixed income, 227
- flow to equity, 369
- follow-on offering, 329
- FOMC, 297
- forecast drivers, 147
- forward dividend yield, 356
- forward earnings, 358
- forward rate, 300
- free cash flow, 140
 - incremental, 150
 - levered, 140, 369
 - to equity, 369
 - unlevered, 140
- free cash flow problem, 532
- free-rider problems, 236
- frictionless markets, 459
- FSA, *see* financial statement analysis
- fundamental value relation, 7
- future value, 9
- general partnership, 575
- geometric series, 51
- goodwill, 631
- Gordon growth model, 346
- GP, *see* general partnership
- gross profit, 628
- gross return, 253
- haircut, 278
- hedge asset, 449

- high minus low portfolio, 486
- HML, *see* high minus low portfolio
- holding period return, 253
- hurdle rate, 135
- hyperinflation, 34
- idiosyncratic risk, 408
- illiquid, 81, 248
- immunization, 286
- impairment charge, 631
- incentive conflict, 323
- income
 - net, 628
 - operating, 628
- income statement, 621
- incremental cash flows, 47, 149
- individual retirement account (IRA), 608
- inflation, 21, 29
 - hyperinflation, 34
- inflation risk, 292
- information asymmetry, 522
- initial public offering, 329, 521
- intangible assets, 631
 - in free cash flow, 144
- interest
 - compound, 70
- interest coverage ratio, 649
- interest rate
 - annual percentage rate (APR), 69
 - annual percentage yield (APY), 73
 - continuous compounding, 109
 - effective annual rate, 70
 - expectations hypothesis, 296
 - forward, 245, 300
 - loan, 67
 - periodic, 69
 - simple interest, 73
 - spot, 245
- interest rate risk, 279
- Interest STRIPS, 238
- internal rate of return, 134, *see* multiple IRRs, 661
 - decision criteria, 134
 - sign rule, 135
- intrinsic value, 591
- inventory, 630
- inverted yield curve, 296
- investment opportunity set, 418
- IPO, 329, 521
 - see* initial public offering, 329
- IRA, *see* Individual Retirement Account
- IRR, *see* internal rate of return, 661
- Jensen's alpha, 484
- key performance indicators, 158
- KPI, *see* key performance indicator
- lease, 100, 553
 - bargain purchase option, 555
 - direct, 554
 - early buyout, 555
 - fixed-price buyout, 555
 - full-service, 554
 - lease-end buyout, 555
 - lessee, 553
 - lessor, 553
 - leveraged-lease, 554
 - market value buyout, 555
 - net, 554
 - rental, 554
 - residual buyout, 555
 - sale-leaseback, 554
 - sale-type, 554
 - security interests, 557

- true, 557
- lenders, 125
- lessee, 100
- lessor, 100
- leverage
 - position, 399
- leverage ratio, 648
 - net, 650
- leveraged recap, 516
- leveraged recapitalization, 516
- levered cost of capital, 126, 504
- levered return, 126, 342, 521
- levered value of the firm, 544
- liabilities, 629
- liability
 - current, 631
 - limited, 633
 - long-term, 632
- liability driven investment, 247
- liability driven investment
 - LDI, 247
- lifetime value, 207
 - LTV-to-CAC ratio, 208
- limit orders, 331
- limited liability, 510, 633
- limited liability company, 575
- limited liability partnership, 575
- limited partnership, 575
- line item, 621
- linear interpolation, 166
- liquid, 236
- liquidating dividends, 346
- liquidation value, 148
- liquidity, 242
 - corporate, 644
- LLC, *see* limited liability company
- LLP, *see* limited liability partnership
- loan, 66
 - borrower, 66
 - collateral, 68
 - covenants, 68
 - creditor, 66
 - lender, 66
 - no interest, 75
 - principal, 67
 - recourse, 508
 - secured, 68
- long-term assets, 630
- long-term value, 577
- loss given default, 289
- LP, *see* limited partnership
- LTV, *see* lifetime value
- LTV-to-CAC, 208
- M&M, 513
- M&M Propositions 1 and 2, 513
- Macauley duration, 283
- margin, 640
 - profit, 640
- margin call, 278
- marginal tax rate, 142
- market cap, 327
- market capitalization, 327
- market efficiency, 258
- market multiple, 360, 592
- market multiples, 574, 591
- market neutral, 475
- market orders, 330
- market portfolio, 459
- market risk, 408
- market risk factor, 462
- market risk premium, 462, 482
- market value balance sheet, 498

- marketable, [232](#)
- marketable securities, [630](#)
- marketed deal commitment, [329](#)
- mean-variance frontier, [410](#)
- method of comparables, [574](#)
- minimum variance portfolio, [413](#)
- MIRR, *see* modified internal rate of return
- mispriced, [276](#), [352](#)
- Modern Portfolio Theory, [391](#)
- modified duration, [284](#)
- Modified Internal Rate of Return, [663](#)
- modified internal rate of return, [195](#)
 - reinvestment rate, [212](#)
- MOIC, *see* multiple on invested cash
- monetary policy, [297](#)
 - monetary easing, [298](#)
 - monetary tightening, [298](#)
- money factor, [102](#)
- money laundering, [129](#)
- monthly recurring revenue, [207](#)
- Moody's, [290](#)
- MPT, *see* Modern Portfolio Theory
- MRR, *see* monthly recurring revenue
- multifactor model, [349](#)
- multiple, *see* market multiple
- multiple on invested cash, [138](#)
- multistage growth model, [363](#)
- MVBS
 - market value balance sheet, [498](#)
- MVP, *see* minimum variance portfolio
- natural hedge, [625](#)
- negative debt, [503](#)
- negative pledge clause, [235](#)
- net debt, [503](#)
- net debt-to-capitalization ratio, [650](#)
- net income, [628](#)
- net leverage, [503](#)
- net long-term investments, [144](#)
- net operating profit after taxes, [141](#)
- net operating profit less adjusted taxes, [141](#)
- net present value, [11](#), [133](#)
- net return, [253](#)
- net working capital, [632](#)
- net working capital investment
 - in free cash flow, [145](#)
- net worth, [629](#)
- NLTI, *see* net long-term investments
- non-cash expense, [626](#)
- non-cash expenses
 - in free cash flow, [142](#)
- Non-diversifiable risk, [408](#)
- non-operating expense, [627](#)
- non-systematic risk, [408](#)
- nontradeable, [232](#)
- NOPAT, *see* net operating profit after taxes
- NOPLAT, *see* net operating profit less adjusted taxes
- NPV, *see* net present value, *see* net present value
- NWCIsenet working capital investment, [145](#)
- OID, *see* original issue discount
- operating earnings, [628](#)
- operating expenses, [624](#)
- operating income, [628](#)
- operating leverage, [625](#)
- opex
 - see* operating expense, [624](#)
- OpEx vs. CapEx, [144](#)

- opportunity cost, 4
 - risk, 5
- order types
 - execution risk, 330
 - price risk, 330
- orders
 - limit, 331
 - market, 330
 - stop, 331
- original issue discount
 - bonds, 303
- OTC, *see* over-the-counter
- over-the-counter, 126
- overhead, 147
 - expenses, 150
- P-STRIPS, 238
- P/E ratio, 359
- P&L, *see* profit and loss statement
- par
 - bonds, 233
- partial sum, 51
- partnership, 575
- payables, 631
- payback period, 136
 - decision criteria, 136
 - discounted, 137
- payment default, 236, 287
- payout ratio, 348
- pecking order, 522
- peg ratio, 362
- pension, 612
- perfect markets, 459
- perpetuity, 27
- PI, *see* profitability index
- plant, property, and equipment, 630
 - gross, 631
 - in free cash flow, 144
 - net, 631
- plowback ratio, 348
- point, 90
- points, 81
- poison put, 235
- portfolio, 262
 - efficient, 412
 - expected returns, 401
 - inefficient, 411
 - market, 459
 - rebalancing, 400
 - returns, 392
 - value weights, 399
 - weights, 394
- portfolio allocation, 395
- portfolios
 - rebalance, 395
- portfolio
 - equal weights, 400
- positive semi-definite, 439
- PP&E, *see* plant, property, and equipment, *see* plant, property, and equipment
- pre-tax expected return, 36
- present value, 6
- present value of growth opportunities, 358
- price return, 333
- price-to-earnings ratio, 359
- primary market, 125
 - stock, 329
- prime borrowers, 86
- Principal STRIPS, 238
- priority of claims, 325
- private equity, 343
- private mortgage insurance, 81

- pro forma
 - P&L, 148
- probability distribution, 288
- profit
 - gross, 628
- profit and loss statement, 621
- profitability index, 198
 - decision rule, 200
 - headcount constraint, 198
- prospective dividend yield, 356
- pure discount bonds, 241
- putable bonds, 239
- PVGO, *see* present value of growth opportunities

- quick ratio, 644

- random variable, 288
- rate risk, 279
- real estate investment trust, 326
- realized returns, 253
- receipts, 623
- receivables, 630
- recurring revenue, 207
- regression line, 403
- reinvestment risk, 272
- REIT, *see* real estate investment trust, 326
- relative mispricing, 276
- relative value, 591
- replicating portfolio, 262
- required cash, 500
- required yield, 243
- residual claimant, 325, 507
- residual value, 101
- retained earnings, 633
- retention rate, 348

- retirement savings accounts, 610
- return
 - abnormal, 484
 - excess, 424
 - excess market, 462
 - nominal, 29
 - real, 29
 - realized
 - stock, 332
- return on assets, 637
- return on equity, 638
- return on invested capital, 638
- return on investment, 134, 638
- returns
 - annualizing, 255
- revenue, 623
- revolving credit, 105
- risk, 5
 - diversifiable, 469
 - exposure, 159, 163
 - firm-specific, 469
 - idiosyncratic, 469
 - market, 469
 - non-diversifiable, 469
 - priced, 469
 - systematic, 469
- risk averse, 434
- risk budget, 449
- risk premium, 289, 424, 461
 - market, 462
- risk tolerance, 435
- risk-free asset, 236, 420
- ROA, *see* return on assets
- ROE, *see* return on equity
- ROI, *see* return on investment, *see* return on investment

- ROIC, *see* return on invested capital
- ROTH retirement savings account, 611
- rule of 40, 653
- S-corp, *see* S-corporation
- S-corporation, 574
- S-corporations, 575
- sales, 623
- cannibalization, 149
 - in free cash flow, 141
 - recognized, 623
 - recurring revenue, 207
- salvage value, 147, 626
- savings account, 227, 228
- high yield, 228
 - traditional, 228
- seasoned equity offering, 329, 521
- secondary market, 125
- stock, 329
- secured bonds, 234
- securitization, 292
- asset-backed, 292
 - collateralized loan obligation, 292
 - pool of loans, 292
- security market line, 480
- selling, general, and administrative
- expense, 625
- sensitivity analysis, 25
- SEO, 329, 521
- SG&A, *see* selling, general, and administrative expense
- shareholder equity, 633
- shareholders, 124, 322
- shareholders equity, 629
- Sharpe ratio, 424
- short position, 398
- short sale, 276
- short squeeze, 278
- short-selling, 11
- short-term investments, 630
- short-term value, 576
- simple interest, 69, 73
- size factor, 485
- small minus big portfolio, 485
- SMB, *see* small minus big portfolio
- SML, *see* security market line
- sole proprietorship, 574
- SPE, 554
- special purpose entities, 554
- special purpose entity
- SPE, 554
- spot interest rates, 245
- spot rates, 245
- spot yields, 245
- Standard & Poors, 290
- statement of cash flows, 634
- statement of income, 621
- statement of operations, 621
- statistical expectation, 258, 288
- stead-state growth, 370
- steady-state, 576, 577
- stock, 322
- adjustable rate preferred, 326
 - closing price, 332
 - common, 325
 - convertible preferred, 326
 - dual class, 325
 - exchange, 327
 - order type, 330
 - par value, 633
 - preferred, 326
 - primary market, 329
 - reverse split, 337

- secondary market, [329](#)
- special dividend, [351](#)
- split, [336](#)
- stated value, [633](#)
- tri-class, [325](#)
- stock buyback, [323](#)
- stock repurchase, [323](#)
- stock vs. flow variable, [156](#)
- stocks
 - Constant Dividend Growth Model, [346](#)
 - Gordon Growth Model, [346](#)
 - liquid, [248](#)
- stop orders, [331](#)
- subprime borrowers, [86](#)
- sunk cost, [151](#)
 - fallacy, [151](#)
- systematic risk, [408](#)
- T-bills, [237](#)
- T-bonds, [238](#)
- T-notes, [237](#)
- tangency portfolio, [427](#)
- tangible assets, [630](#)
- target, [247](#)
- tax
 - capital gains, [153](#)
- tax basis, [153](#)
- tax rate
 - average, [142](#)
 - effective, [142](#)
 - marginal, [142](#)
- tax shield
 - depreciation and amortization, [144](#)
- taxes
 - capital gains, [374](#)
 - coupon bonds, [304](#)
- technical default, [236](#), [287](#), [532](#)
- tender, [519](#)
- terminal value, [363](#), [370](#), [577](#)
- terminal year, [577](#)
- theoretical spot rate curve, [266](#)
- time deposits, [231](#)
- time diversification, [432](#)
- time value of money, [10](#)
- timeline, [2](#)
- TIPS, *see* Treasury Inflation Protected Securities [239](#), *see* treasury
- total net payout, [368](#)
- Total Payout Model, [366](#)
- total payout ratio, [366](#), [367](#)
- total return, [333](#)
 - bonds, [272](#)
- trade cash cycle, [647](#)
- trade credit, [632](#)
- trailing twelve months, [360](#), [599](#)
- transaction cost, [242](#)
- treasury
 - inflation protected securities (TIPS), [293](#)
- Treasury bills, [237](#)
- Treasury bonds, [238](#)
- Treasury Inflation Protected Securities, [239](#)
- Treasury notes, [237](#)
- Treasury securities, [236](#)
- Treasurys, [236](#)
- tri-class shares, [325](#)
- TTM, *see* trailing twelve months, *see* trailing twelve months
- turnover, [623](#)
- UE, *see* unlevered earnings
- underinvestment, [533](#)

- underwrite, [327](#)
- UNI, *see* unlevered net income
- unique risk, [408](#)
- unlevered cost of capital, [511](#), [544](#)
- unlevered earnings, [141](#)
- unlevered free cash flow, [141](#)
- unlevered net income, [141](#)
- unlevered value of the firm, [544](#)
- usable life
 - asset, [147](#)
- valuation multiple, [360](#)
- value, [1](#)
 - Costs and benefits, [1](#)
 - creation, [11](#)
- value-growth factor, [485](#)
- variable cost, [624](#)
- venture capital, [326](#)
- volatility, [339](#)
- voting rights, [325](#)
- WACC
 - pre-tax, [502](#)
- WACC valuation, [548](#)
- weighed average cost of capital
 - pre-tax, [502](#)
- working capital, [632](#)
 - liquidity, [632](#)
- write down, [631](#)
- yield
 - bond equivalent, [243](#)
- yield curve, [244](#)
- yield-to-maturity, [243](#)
- YTM, [243](#)
- zero coupon bonds, [233](#)
- zero-beta assets, [469](#)
- zeros, [233](#)