Financial Decision-Making

Michael R. Roberts
The Wharton School, University of Pennsylvania

Current Draft Date: April 3, 2023

Copyright © 2020 Michael R. Roberts

Contents

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

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. Any references to third parties are for demonstrative or descriptive purposes only.

Copyright © 2020 Michael R. Roberts

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.

The problem is that few of us have the time, money, and patience to learn finance from traditional textbooks otherwise known as "door stoppers" - books so long and heavy they're more likely used as a door stop than as a learning instrument. So, I wrote this book based on my lectures to Duke and Wharton undergraduates and MBAs, as well as executives, that I have had the privilege of teaching over the last 22 years. It is meant to be succinct, engaging, informal (but precise), emphasizing intuition and applications. The goal of this book 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 found textbooks are often written more for the academics selecting them than the students using them.) This required making difficult choices. Some topics are lightly covered (e.g., capital markets, working capital management) or excluded (e.g., derivatives, mergers & acquisitions). Proofs of 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 interested in more detail or topics not covered here, there are countless other resources with which you'll be well prepared to engage after reading this book.

Instead, I decided to emphasize real-world applications to illustrate the thought process

Copyright © 2020 Michael R. Roberts

for tackling financial challenges that most of us face in our lives. This makes the book more challenging, but hopefully more valuable. To ease the cognitive burden, I've 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 - corporate finance, personal finance, investments, etc. This emphasis also ensures readers are able to adapt to an ever-changing financial environment.

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,
- business cases that mimic real life scenarios in which decision-makers find themselves (for instructors),
- end-of-chapter exercise solutions and business case teaching notes (for instructors), and
- detailed lesson plans and activity suggestions (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 - saving for retirement; going to college; financing a home, car, school, etc. Part II: Decisions Most People Make, examines financial decision-making in business settings and details different investments - bonds, stocks, and portfolios. This part complements Part I, as for most people it is important to understand the basics of business decision-making - which are no different from personal decision-making - and where to save one's money. Part III: Decisions Finance Specialists 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.

Copyright © 2020 Michael R. Roberts

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.

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

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, Gumby, 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, Daniel McFadden, Andrew Metrick, Angela Merrill, Roni Michaely, Mitchell Petersen, James Powell, Manju Puri, Matt Richardson, Nick Roussanov, Mike Schwert, Rob Stambuagh, Amir Sufi, Vish Viswanathan, Jessica Wachter, Toni Whited, Robert Whitelaw, Amir Yaron, Moto Yogo, Rebecca Zarutskie, and Jaime Zender.

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.

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Part I

Decisions Everyone Makes

Copyright © 2020 Michael R. Roberts

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 (addition, subtraction, multiplication, and division).

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

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

Copyright © 2020 Michael R. Roberts

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, this process of weighing costs and benefits is done automatically. Nonetheless, it is illustrative. We make decisions by weighing the costs and benefits whose relative sizes determine the value of the decision to us.

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 these ideas 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.

Let's first rewrite this financial value equation to recognize *when* we receive and pay money. Figure 1.3 illustrates the timing of cash flows.

Copyright © 2020 Michael R. Roberts

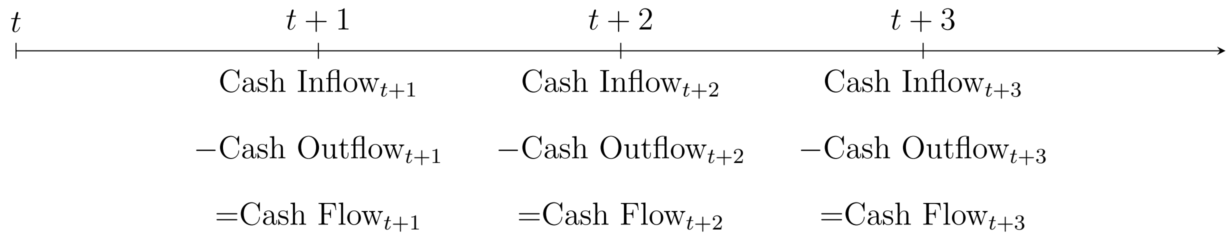

Figure 1.3: Timeline and Cash Flows

The figure presents a **timeline** showing cash inflows and outflows by time period, which is indicated by the labels on the 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.

Now we can express financial value as

$$Value_t = CashFlow_{t+1} + CashFlow_{t+2} + CashFlow_{t+3} + ... \tag{1.1}$$

The equation says that the value of a decision at time t is equal to the sum of all the money we will receive in the future less any payments we have to make along the way. The ellipsis ("...") at the end of the equation indicate that the cash flows could go on indefinitely as suggested 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.

Copyright © 2020 Michael R. Roberts

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. (In case you're wondering, opportunity cost is different from inflation, which we'll discuss later.) 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, what else would we do with the money and what would that money 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 2017, 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

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

Copyright © 2020 Michael R. Roberts

movements or jaggedness of each line. The blue line corresponding to the small-cap stocks' 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!

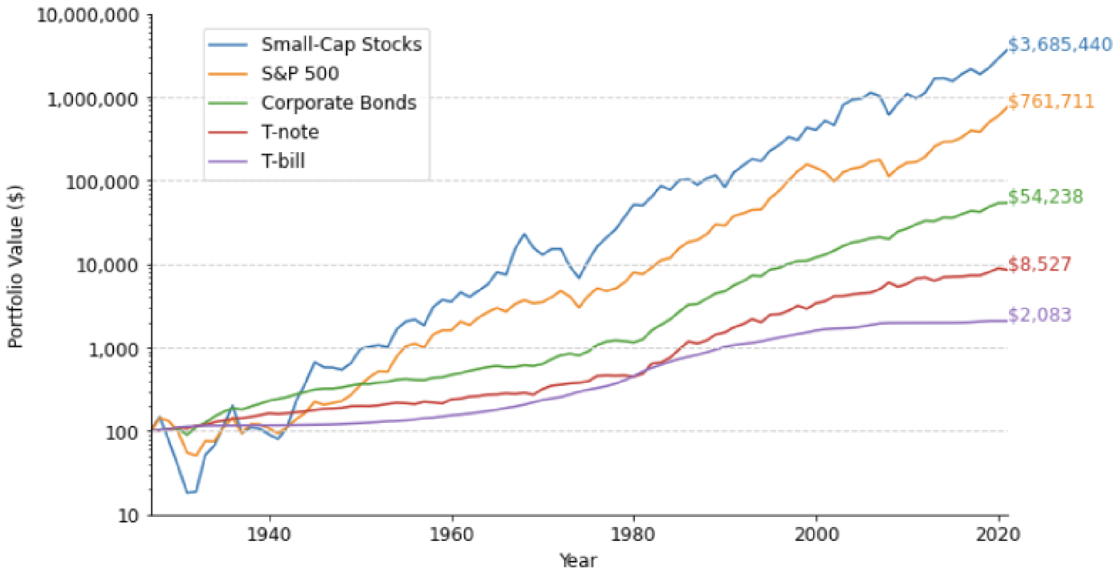

Figure 1.5: Value of \$100 Investments from 1927 to 2021

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.

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

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

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 a lot of things depending on the context: **cost of capital**, **discount rate**, and **expected return**. We'll see even more names later. What's important to understand is that r reflects the risk of the cash flows - riskier cash flows, higher r.

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

[2]Quick math refresher: raising quantities to a power is just shorthand for multiplication. For example, $2^4 = 2 \times 2 \times 2 \times 2 = 16$. So, $(1 + r)^3 = (1 + r) \times (1 + r) \times (1 + r)$.

Copyright © 2020 Michael R. Roberts

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 compute it.**

Equation 1.2 shows that anything affecting (i) future cash flows or (ii) opportunity costs 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 a financially beneficial decision. 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 $100 million of 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 discount rate.

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 therefore

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

We should pay no more than $179.71 in exchange for receiving $100 each year for the next two years when the opportunity cost of our money is 7.19%.

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.

1.4 Price vs. Value

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

Copyright © 2020 Michael R. Roberts

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 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 for each share we purchase today.

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.

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

We've made two changes to equation 1.2. We subtracted 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. As with the term 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 \$179.71. The net present value of the deal when the price is \$200 is $-200 + 179.71 = -\$20.29$, which is value-destructive for us. Paying \$179.71 for the deal means the net present value is $-179.71 + 179.71 = 0$. 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 $-160 + 179.71 = \$19.71$, 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 in certain circumstances when the price is greater than 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 and are no better or worse off.

Copyright © 2020 Michael R. Roberts

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, even though there is a financial cost to 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 in this book. 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 and what kind of return we will earn when we sell 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 as well! The point is that financial value is often one part of a broader trade-off when we make decisions, albeit an important part.

That said, there are times when nonfinancial considerations can lead us astray in what are primarily financial decisions. A good example is paying off low interest rate debt before high interest rate debt. People may feel better about having paid some of their debt, but they are unambiguously making themselves financially worse off and creating other problems they may not appreciate. Financial proficiency can help mitigate psychological biases that lead us to make bad financial decisions.

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, Wharton offers over 30 different finance courses that can appear completely unrelated. Some course titles include: Corporate Valuation, Venture

Copyright © 2020 Michael R. Roberts

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 an exotic 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 is 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 identifying the cash flows and discount rates in different applications. This is on what we'll spend our time.

1.6 Key Ideas

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

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

  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). This is called a **discounted cash flow** or **DCF** model.

- 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.

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

Copyright © 2020 Michael R. Roberts

- The opportunity cost (a.k.a., discount rate, expected return, cost of capital) r, reflects the risk of the cash flows. The greater the risk, the greater the opportunity rate.

- The expression "a dollar today is worth more than a dollar tomorrow" comes from the concept of opportunity cost and follows immediately from our value equation. Dividing cash flows by $(1+r)^t$ makes them smaller because we are dividing by a number larger than one (e.g., 1 + .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. We'll explore this intuition more deeply in the next chapter.

- Financial value is forward-looking, based entirely on what is *going to* happen, not what *has* happened. In fact, considering money spent in the past when making a decision is a common error known as **the sunk cost fallacy**. 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.

Copyright © 2020 Michael R. Roberts

Chapter 2

Retirement Savings and the Value of College

Fundamental value relation

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

This chapter

- provides further intuition for the fundamental value relation,
- 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, and if so, should we finance it or pay for it from savings?

It's actually surprising (to me) that with nothing more than what we'll learn in the first two chapters, we can tackle these important questions.

Copyright © 2020 Michael R. Roberts

2.1 Intuition for Opportunity Cost Adjustment

Our fundamental value relation (equation 1.2) says that the value of a decision or an asset at any point in time, t, is equal to the sum of the future costs and benefits, as measured by cash flows, adjusted for their opportunity cost, as measured by r. While the first part of value measurement - the sum of future costs and benefits - is fairly intuitive, the second part - adjusting for opportunity cost - is less so. Let's understand why we divide the cash flows by $(1+r)^t$. To do so, a couple of analogies are useful.

Imagine there are two boxes. (Bear with me.) One box weighs 125 pounds (lb), 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 boxes' weights 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.

$$\text{Weight in kilograms:} \quad 125 \text{ lb} \times \tfrac{1 \text{ kg}}{2.205 \text{ lb}} + 59 \text{ kg} = 115.69 \text{ kg}$$

$$\text{Weight in pounds:} \quad 125 \text{ lb} + 59 \text{ kg} \times \tfrac{2.205 \text{ lb}}{1 \text{ kg}} = 255.10 \text{ lb}$$

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

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**.

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.

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

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

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

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.

Copyright © 2020 Michael R. Roberts

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 $1/(1+r)^s$ for $s = t+1, t+2, t+3, ...$ act as exchange rates for time, converting the time unit of each future cash flow into period t time unit. Let's see this conversion process in action with an application.

2.2 Application: Cost of College

2.2.1 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**. A timeline in finance is nothing more than a horizontal line with markings indicating different time periods and cash flows. It is a simple and surprisingly useful visual aid for solving financial problems.

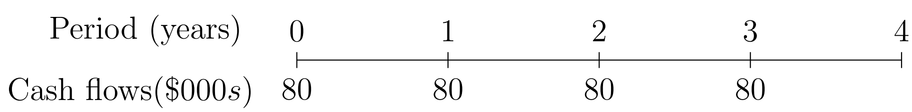

Figure 2.1: School Savings Timeline

The timeline in figure 2.1 shows our annual payments of $80,000 per year for four years beginning at the start of each year. Notice that today, the time of the decision, is indicated by period 0. This is a common convention that eases the mapping of the timeline to the fundamental value relation. We could have put negative signs in front of each cash flow to indicate cash outflows, but their meaning is unambiguous here.

Ignoring opportunity costs, we would need $320,000 ($4 \times \$80,000$) 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

Copyright © 2020 Michael R. Roberts

money. Because we have to make the payments to complete school, any investment we take should be relatively safe. For example, we shouldn't invest our money in stocks unless we are willing to risk not being able to pay for school. This is an important point mentioned previously but worth repeating. The opportunity cost for a set of cash flows must reflect the risk of those cash flows - higher risk, higher opportunity cost. How to measure that risk and estimate the opportunity cost comes later.

Assume we invest our money in a bank savings account - very safe - earning 5% per year, i.e. $r = 0.05$. The value of our cash flows in *today's* dollars is

$$Value_0 = 80,000 + \frac{80,000}{(1+0.05)} + \frac{80,000}{(1+0.05)^2} + \frac{80,000}{(1+0.05)^3} = \$297,859.84.$$

We don't adjust the first payment because it happens "now," in period 0, at the start of school. Also, note that we need substantially less than $320,000 to pay for all four years of school because any money we have at the start of school will earn interest and grow over time.

A useful way to visualize the calculations and relate them to time unit conversions is provided in figure 2.2. Each cash flow has a different time unit corresponding to the period in which it is made. The first payment made at the start of school has a time unit of 0, the second a time unit of 1, the third a time unit of 2, and the fourth a time unit of 3. To add these payments together, we need to convert them to a common time unit. Because we want to know how much money we'll need at the start of school, we convert them to time unit 0.

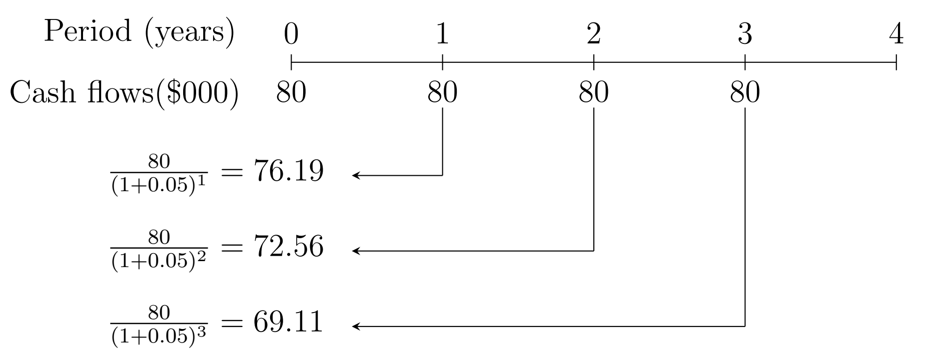

Figure 2.2: Time Unit Conversion - Moving Money Back in Time

The first payment made at the start of school already has time unit 0 so no conversion is needed. The second payment made one year later has a time unit of 1 and can be converted to time unit 0 by dividing by $(1 + 0.05)^1 = 1.05$. The $76,190 in figure 2.2 is the **present value** of $80,000 one year later at an opportunity cost of 5%. Put differently, we

Copyright © 2020 Michael R. Roberts

are indifferent (i.e., equally happy) between paying \$76,190 today or \$80,000 one year from today because we can invest \$76,190 for one year at 5% and have \$80,000 at the end of the year, $\$76,190 \times (1+0.05) = \$80,000$. Using our box analogy from above, we don't care if we have to carry a box weighing 100 pounds or a different box weighing 45.25 kilograms; both boxes are equally heavy.

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%, namely, \$297,859.84 as computed above. This process of moving cash flows *back* in time by dividing by $(1+r)^t$ is called **discounting**.

Another way to see what's happening in our value calculation is shown in table 1, called an **amortization table** or **amortization schedule**. Today, we deposit (i.e., put) \$297,860 into a bank account earning 5% per year. We immediately withdraw (i.e., take out) \$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. One year later, our savings has earned $0.05 \times 217,859.84 = \$10,892.99$ in interest and we have to withdraw another \$80,000 for the second payment. The account now has $217,859.84 + 10,892.99 - 80,000 = \$148,752.83$ 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				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

2.2.2 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 $(1+r)^s$ to move future cash flows back in time to get

Copyright © 2020 Michael R. Roberts

their present values, now we *multiply* the cash flows by $(1+r)^s$ to move cash flows forward in time to get their **future values**. Figure 2.3 illustrates this process.

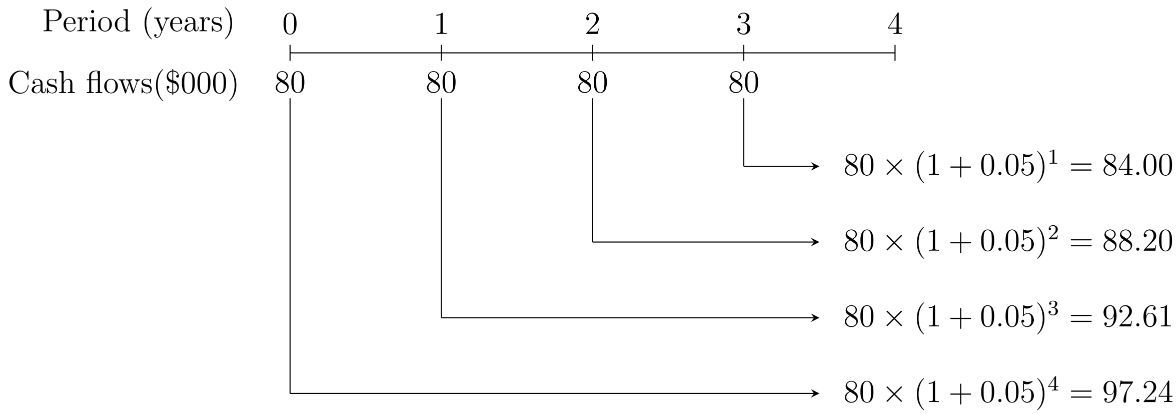

Figure 2.3: Time Unit Conversion - Moving Money Forward in Time

The figure shows how each cash flow's time unit is converted to a time unit 4, 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.

$$Value_4 = 80 \times (1+0.05)^4 + 80 \times (1+0.05)^3 + 80 \times (1+0.05)^2 + 80 \times (1+0.05) = \$362,050.50$$

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

Note, \$362,050.50 is the same number we get if we multiply the present value of the cash flows, \$297,859.84, by $(1+0.05)^4$. This is not a coincidence. We can move sums of money around in time or each individual component of the sum. It doesn't matter.

2.2.3 Summary

We introduced several fundamental concepts in this section worth summarizing before moving on to further applications.

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 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.

Copyright © 2020 Michael R. Roberts

(a) Multiplying a cash flow by the discount factor (i.e., dividing by $(1 + r)^s$) moves money *back* in time by s 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 $(1 + r)^s$) moves money *forward* in time by s periods. The process of moving money forward in time is called **compounding**.

3. **Present value** (PV) is the value of a cash flow as of today (period 0). The present value of a single cash flow T periods in the future is

$$\text{Present value} = \frac{CashFlow}{(1+r)^T} \tag{2.1}$$

4. **Future value** (FV) is the value of a cash flow at a future point in time. The future value of a single cash flow T periods in the future is

$$\text{Future value} = CashFlow \times (1+r)^T \tag{2.2}$$

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, also known as **inflation**. 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

Copyright © 2020 Michael R. Roberts

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 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 price changes 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. The real return is denoted by rr and defined as follows.

$$rr = \frac{1+r}{1+\pi} - 1. \tag{2.3}$$

The nominal return is denoted by r and corresponds to the opportunity cost we've been talking about since the previous chapter. 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 slower or faster than inflation.

An often used approximation for equation 2.3 is

$$rr = r - \pi. \tag{2.4}$$

This approximation works well when nominal rates and inflation are small (less than 10%).

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, our purchasing power is declining by almost 1.87% because of inflation. So, our money increases, but our ability to buy stuff, in this case oranges, decreases. 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.

Copyright © 2020 Michael R. Roberts

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 2.4. 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 2.4: School Costs with Inflation (Nominal Cash Flows)

Real cash flows measure the purchasing power of money relative to a base year. 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 2.5

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

Figure 2.5: 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. Put differently, imagine buying $80,000 of "stuff" - goods and services - today. To buy that same stuff one year from today, we would need $81,600. To buy that same stuff two years from today, we would need $83,232. And, so on. So, the $80,000 real cash flow each year is telling us the corresponding purchasing power of the nominal cash flow each year.

Now assume that the cost of college increases by 4% per year. The timeline with nominal cash flows is shown in figure 2.6.

Again, the nominal cash flows describe the actual money we have to pay each year to go to college. To determine the corresponding real cash flows relative to today, we deflate (divide

Copyright © 2020 Michael R. Roberts

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 2.6: School Costs Outpacing Inflation (Nominal Cash Flows)

by $(1+\pi)^t$)) each nominal cash flow. Figure 2.7 shows the timeline with the corresponding real cash flows.

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 2.7: School Costs Outpacing Inflation (Real Cash Flows)

Notice now that the real cash flows are all greater than $80,000. The cost of college is increasing more rapidly than inflation or prices in general.

Taking a step back, we can see how high inflation can be problematic. If the price 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. This situation is exactly what occurred in 2021 and 2022 in the U.S. when inflation shot up to over 9%, while interest rates and wage growth were relatively low - approximately 4% per year for each. 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 valuation 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. Specifically, **we discount nominal cash flows by nominal returns, or real cash flows by real returns.**

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 2.6. With a

Copyright © 2020 Michael R. Roberts

nominal discount rate of 5%, the present value of the cost of college is

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

The real cash flows for college costs are shown in figure 2.7. 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

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

We get the same result.[1]

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 Application: 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. To clarify the process, we'll keep things as simple as possible by making some unrealistic assumptions that we'll relax in the next section.

2.4.1 Timeline

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 is expected to live to 87. Figure 2.8 shows the start of a timeline based on this information.

Period 0 as always corresponds to today, when Sophie is 30. She will work for 40 years and then retire at age 70. While working, she'll save money so that during retirement she can

[1] Mathematically, what's happening is that the formulation using real values is equivalent to dividing the numerator and denominator of each term by $(1+\pi)^t$.

Copyright © 2020 Michael R. Roberts

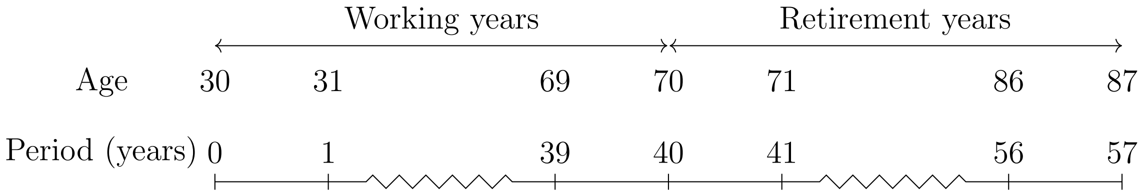

Figure 2.8: Retirement Savings Timeline

use those savings to take care of her needs. A savings retirement plan consists of forecasts of her consumption needs in retirement and savings strategy while working. In other words, we need to populate our timeline with cash flows.

Before doing so, it's important to 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? These are questions we will explore after we've got a baseline retirement plan in place.

2.4.2 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 housing, food, and healthcare; 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. Nonetheless, one starting point for this exercise is Sophie's current expenditures grossed up by inflation. Assuming Sophie spends $61,311.37 per year today, her needs in retirement 40 years from now at an expected inflation rate of 3% per year are $61,311.37 \times (1 + 0.03)^{40} = \$200,000$.

The obvious shortcoming of this approach to determining Sophie's monetary needs in retirement is that her consumption and expenditures will be very different 40 years from today. For example, medical expenses are likely to become much more important in retirement. Her tastes and lifestyle will likely change, perhaps becoming more expensive. She may no longer have mortgage or rent payments. And so on. However, to keep the focus on the mechanics, we will simply assume that Sophie needs $200,000 per year in retirement to live the lifestyle she wants until she passes. These cash flows are illustrated in Figure 2.9.

Copyright © 2020 Michael R. Roberts

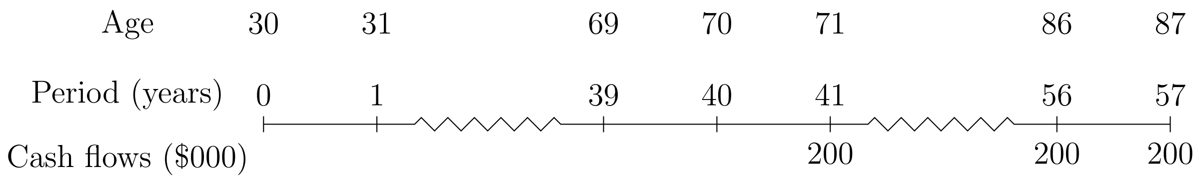

Figure 2.9: Retirement Savings Timeline - Retirement Cash Flows

2.4.3 The Nest Egg

Question two: 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.

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).

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

Sophie needs approximately \$2.63 million at the start of her retirement. These savings for retirement are sometimes referred to as a **nest egg**, whose calculation is visualized in Figure 2.10

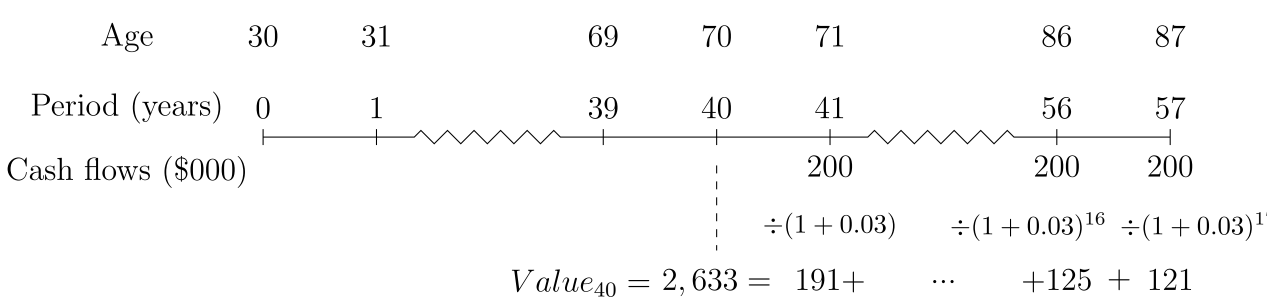

Figure 2.10: Retirement Savings Timeline - Nest egg

Copyright © 2020 Michael R. Roberts

Sophie's retirement needs - $200,000 a year for 17 years - are called an **annuity**, which is defined as 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.

$$\text{Value of an Annuity}_t = \frac{CF}{r} \times \left(1 - (1+r)^{-(T-t)}\right) \tag{2.5}$$

Recall that CF stands for $CashFlow$, r is the discount rate, and T is time of the last cash flow or number of cash flows. (See the Technical Appendix for a proof of this result.) Applying this result to Sophie's retirement needs yields

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

the same amount we computed above using our fundamental value relation.

A few comments regarding equation 2.5 are in order. First, it 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. Second, it 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 result can be used to answer many other questions not so easily answered by brute force calculations in a spreadsheet.

Finally, when using equation (2.5) to find the value of an annuity, the first cash flow is assumed to arrive 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,

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

exactly what we computed above.

2.4.4 The Savings Strategy

Question three: How much does Sophie need to save each year during her working years to ensure she has $2.633 million at the start of her retirement? Let's start by finding the

Copyright © 2020 Michael R. Roberts

present value of this amount by discounting from period 40 to period 0 (i.e., today) using the same discount rate we use in retirement, 3%.

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

Figure 2.11 illustrates how this calculation converts the $2.633 million forty years from today into $807,232 in today's dollars (i.e., present value).

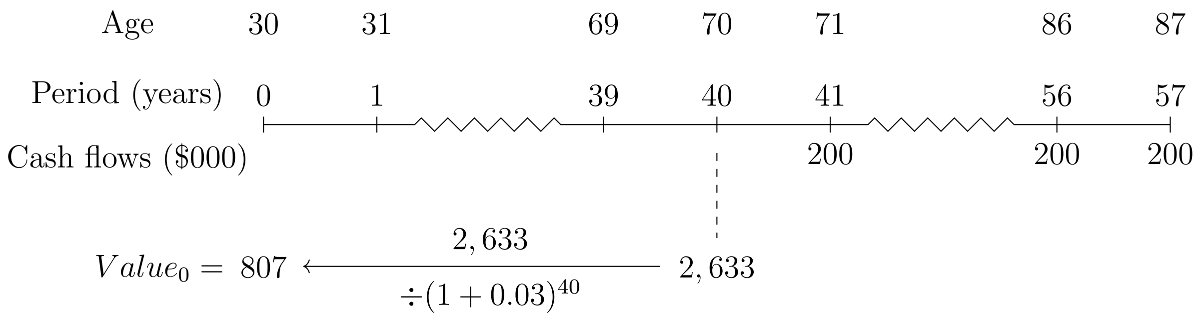

Figure 2.11: Retirement Savings Timeline - Present Value of Nest egg

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.

Let's assume Sophie doesn't have any savings today and wants to contribute a constant amount each year to a savings account paying 3%. How can we figure out how much she needs to contribute each year to reach her goal? This is where the annuity result (equation 2.5) comes in handy. Rather than plugging in the cash flow and discount rate to get the present value, we have the present value of Sophie's savings annuity, $807,232, and the discount rate, 3%. We just need to solve for the annuity cash flow.

A little algebra shows that equation 2.5 can be rewritten as

$$CF = \frac{\text{Value of Annuity}_t \times r}{1-(1+r)^{-(T-t)}} \tag{2.6}$$

Plugging into the equation the present value of Sophie's nest egg, her investment return during her working years, and the length of her working years produces

$$CF = \frac{807,232 \times 0.03}{1-(1+0.03)^{-(40-0)}} = \$34,922.81. \tag{2.7}$$

Copyright © 2020 Michael R. Roberts

With no current savings, Sophie needs to save $34,922.81 every year for 40 years starting one year from today while earning 3% per year in interest. The completed timeline and retirement savings strategy is illustrated in Figure 2.12.

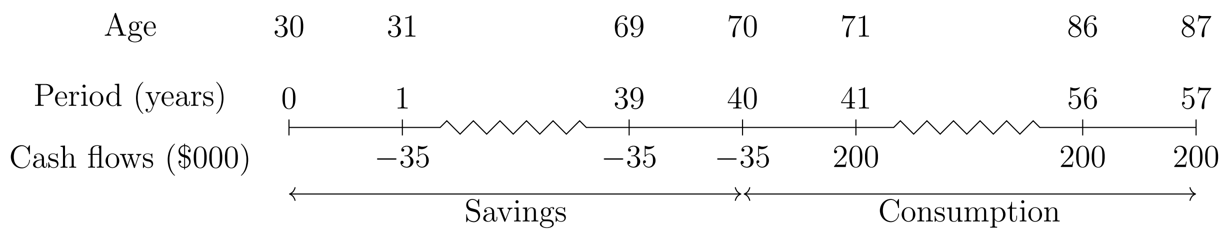

Figure 2.12: Retirement Savings Strategy

2.4.5 What if...?

With any financial model, and that's exactly what we've constructed, we should always perform sensitivity or robustness analysis. That's just fancy-speak for asking "What if?" of our model. Below are some examples.

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 had to save to reach her nest egg goal. In other words, we can deduct this money from the present value of her nest egg, $807,232 - 250,000 = \$557,232$. This calculation is sensible because both values have time unit 0. Using this adjusted present value, we can compute her new annual savings.

   $$\text{CF} = \frac{557,232 \times 0.03}{1 - (1 + 0.03)^{-40}} = \$24,107.21$$

   With $250,000 of savings at age 30, Sophie only needs to save an additional $24,107.21 per year (compared to $34,922.81) earning 3% per annum to reach her nest egg of $2.633 million by age 70. (Of course, her $250,000 needs to be saved and earn 3% per year as well.)

2. What if Sophie wants to leave an inheritance (i.e., bequest) of $500,000 for her kids when she dies? The cash flows in retirement are now illustrated in figure 2.13.

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

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

Copyright © 2020 Michael R. Roberts

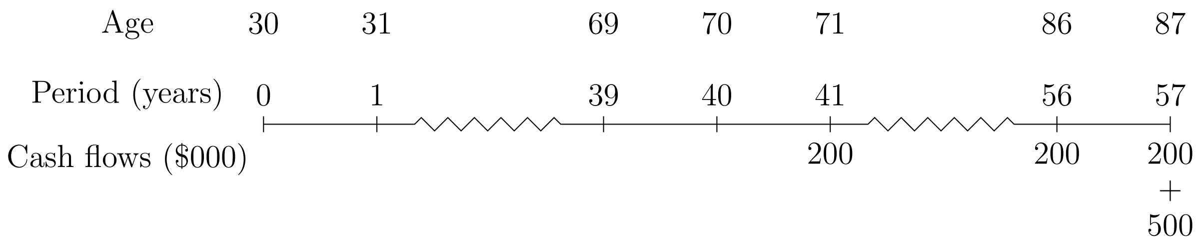

Figure 2.13: Retirement Savings with a Bequest

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

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

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 is illustrated in figure 2.14. Note there is no age after 87 because Sophie is assuming she will have passed on at 87.

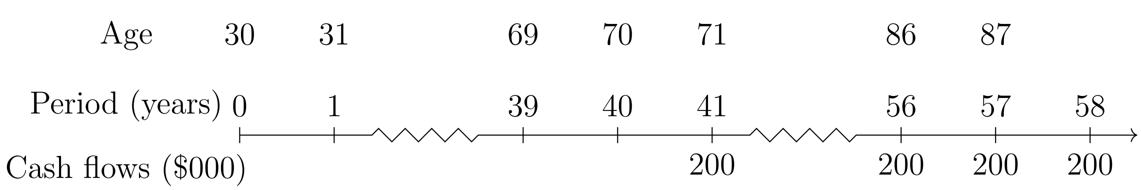

Figure 2.14: Retirement Savings with a Perpetual Bequest

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

Copyright © 2020 Michael R. Roberts

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, she will have to leave her descendants a lump sum of

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

This amount ensures that her descendants can withdraw \$200,000 per year forever.

What we just did is compute the present value of a **perpetuity**, which is just like an annuity except the cash flows go on forever ($T = \infty$). The general relation for computing the present value of a perpetuity is

$$Value_t = \frac{CF}{r} \tag{2.8}$$

4. *Challenging:* If Sophie decides to consume less in retirement, say \$175,000 per year, how much longer will her nest egg last? Remember, Sophie's consumption in retirement is just an annuity whose present value is given by equation 2.5. If we solve that equation for T, the length of the annuity, we get

$$T = -\frac{\ln\left(1 - \frac{Value_0 \times r}{CF}\right)}{\ln(1+r)} \tag{2.9}$$

where ln is the natural logarithm.[2] Now, we can just plug and chug

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

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. Because outliving one's money is a big concern for retirees, it is important to understand just how far we can stretch our money.

There are many other "what if" questions we could ask including

- 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?

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 pass at Sophie's retirement plan explores some of these changes.

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

Copyright © 2020 Michael R. Roberts

2.5 Application: Saving for Retirement (Inflation and Taxes)

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

- needs \$200,000 at the end of her first year in retirement, after which her consumption needs to grow by 2% per year to keep up with expected inflation.
- 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 3%;
- is taxed at a constant rate of 35% on each annual withdrawal during retirement; and,
- will grow her savings each year while working by 4%, her anticipated average annual salary increase.

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

2.5.1 Consumption in Retirement

The timeline illustrating Sophie's retirement consumption is presented in Figure 2.15. Remember, these cash flows represent the money she needs each year in retirement to purchase food, housing, clothes, vacation, medicine, etc. As before, we'll take these cash flows at face value, leaving an investigation into how to actually forecast these cash flows as an exercise.

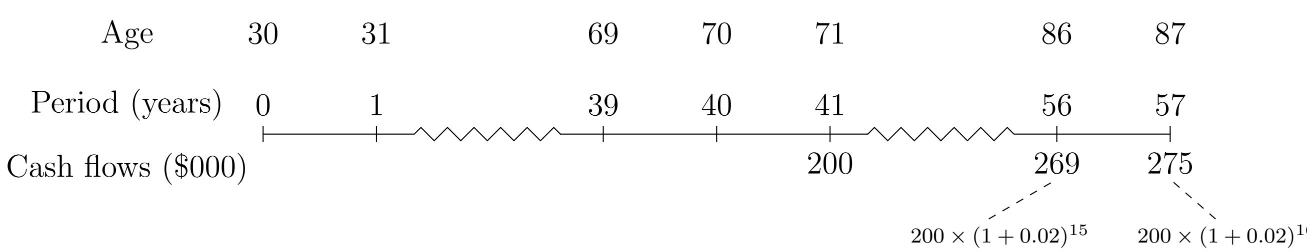

Figure 2.15: Retirement Consumption with Inflation

Copyright © 2020 Michael R. Roberts

2.5.2 The Nest Egg

We can't use our annuity formula (equation 2.5) from earlier because the cash flows are not constant. However, because the cash flows *grow at a constant rate*, we can use a modified version of our annuity formula.

$$Value_t = \frac{CF}{r-g} \times \left(1 - \left(\frac{1+r}{1+g}\right)^{-(T-t)}\right) \tag{2.10}$$

where CF is the cash flow one period ahead of when we are valuing the growing annuity, and g is the constant cash flow growth rate.

Equation 2.10 returns the value of a **growing annuity**. In fact, when the growth rate of the cash flows is zero, equation 2.10 reduces to equation 2.5. No growth is zero growth. The rules for using equation 2.10 are otherwise the same as before - cash flows that are (i) equally-spaced in time and (ii) over a finite horizon - except now the cash flows are allowed to change in a very specific way.

Applying this result to Sophie's retirement needs tells us that the size of her nest egg at the start of retirement (period 40, age 70) should be:

$$Value_{40} = \frac{200,000}{0.03-0.02} \times \left(1 - \left(\frac{1+0.03}{1+0.02}\right)^{-(57-40)}\right) = \$3,056,618.26$$

One caveat with growing annuities is that when the discount rate and the cash flow growth rate are equal ($r = g$), then we can't use equation 2.10. (We can't divide a number by zero.) Instead, we have to use the following equation.

$$Value_t = \frac{CF \times (T-t)}{1+r} \tag{2.11}$$

Taxes

Now, let's bring in some taxes. Each year Sophie withdraws money in retirement, the government takes some in taxes. Exactly how much they take depends on several factors, most importantly the tax policy in place when she retires. In addition, from where the money is coming - a traditional retirement account (e.g., 401-k, IRA), a ROTH retirement account, private savings - and how it is coming - asset sales versus income generated by the assets (dividends, interest) - also matter, at least as of 2022.

Copyright © 2020 Michael R. Roberts

Regardless, the first implication of taxes is that Sophie must withdraw more than what she needs to consume each year so she can pay her taxes. How much more?

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

where τ is her tax rate which we'll assume is 35%. In other words, for each dollar she withdraws from her savings in retirement, she only keeps $(1-\tau) = (1-0.35) = \$0.65$.

Let's modify her nest egg computation to account for taxes. (We'll use nominal cash flows and the nominal discount rate.)

$$Value_{40} = \frac{200,000/(1-0.35)}{0.03-0.02} \times \left(1 - \left(\frac{1+0.03}{1+0.02}\right)^{-17}\right) = \$4,702,489.64$$

Holy Schnike! Our nest egg has to be 1.7 million dollars, or approximately 56%, larger because of the taxes. Now we see why tax planning can be so important, especially for people in higher tax brackets. However, Sophie's more aggressive investment strategy during her working years leads to a present value for her nest egg equal to

$$Value_0 = \frac{4,702,489.64}{(1+0.10)^{40}} = \$103,901.17$$

Figure 2.16 illustrates the effects of inflation and taxes on Sophie's nest egg.

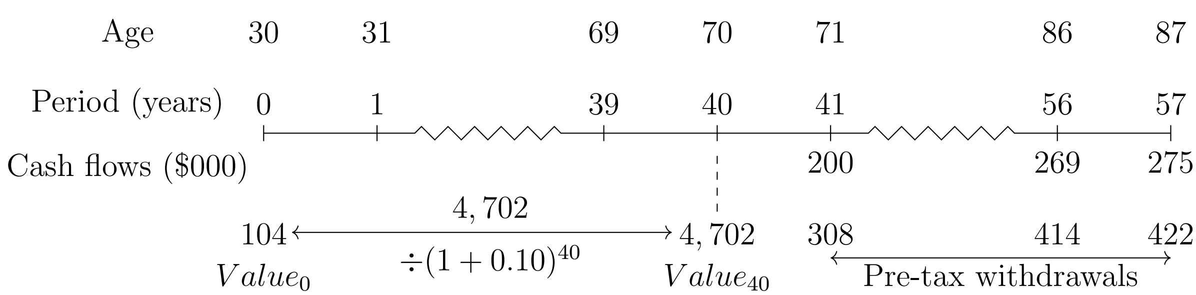

Figure 2.16: Retirement Consumption with Inflation

2.5.3 The Savings Strategy

To compute how much Sophie needs to save each year while working, we can solve the growing annuity formula for the cash flow, CF.

$$CF = \frac{Value_t \times (r-g)}{\left(1 - \left(\frac{1+r}{1+g}\right)^{-(T-t)}\right)} \tag{2.12}$$

Copyright © 2020 Michael R. Roberts

The cash flow from this equation is the first cash flow of the growing annuity, one year from today. 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

$$CF = \frac{103,901.17 \times (0.10 - 0.04)}{\left(1 - \left(\frac{1+0.10}{1+0.04}\right)^{-40}\right)} = \$6,973.84.$$

While this amount seems low, remember that she plans on increasing it by 4% every year for 40 years. Additionally, she expects to earn 10% per year with her more aggressive savings strategy. To find how much she needs to save for any other year, we just need to "grow" this initial savings by the 4%. For example, her last savings just before retiring will be $\$6,973.84 \times (1 + 0.04)^{39} = \$32,193.81$.

Sophie's final savings plan is illustrated in Figure 2.17.

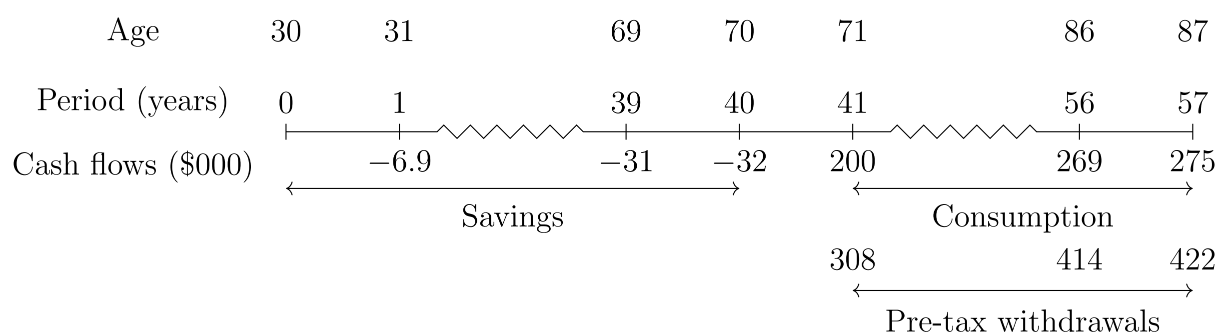

Figure 2.17: Retirement Strategy with Taxes & Inflation

2.5.4 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 forecast 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.

Copyright © 2020 Michael R. Roberts

- "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, etc.) to address potential risks. Doing so is easy. You're just changing numbers. The calculations are the same.

- "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 readers may worry that they don't make enough or can't save now because times are tight. First, saving a little over long periods of time can have a big impact. If you save $1,200 a year for 40 years at an 8% annual return, you'll have $310,867.82 dollars at the end of the 40 years. That's just $100 a month, approximately. Second, not being able to save now or in the near future does not change the importance of the exercise. It only means that our savings for the next few years will be zero.

2.6 Application: Should We Go To College?

Does it make financial sense to go to college? Let's assume the following.

- School costs $80,000 per year payable at the start of each year for four years.

- If we skip school, we will work for 50 years earning $30,000 after tax at the end of the first year. This amount will grow by 2% per year thereafter.

- If we go to school, we will work for 46 years earning $70,000 after tax one year after graduating. This amount will grow by 2% per year thereafter.

- Our opportunity cost of capital is 5% per year.

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.

Copyright © 2020 Michael R. Roberts

2.6.1 Skip College

If we skip college, the only cash flows relevant for valuation are our after-tax earnings, which are presented in figure 2.18.

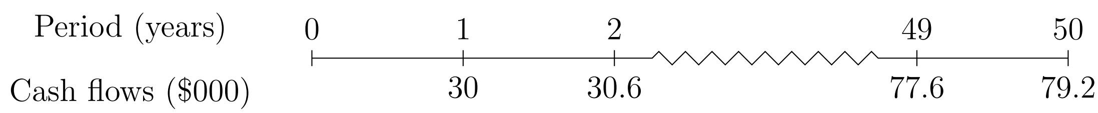

Figure 2.18: Skip College Timeline

We can compute the value of these cash flows by discounting and summing them like so.

$$Value_0 = \frac{30,000}{1+0.05} + \frac{30,600}{(1+0.05)^2} + ... + \frac{79,164}{(1+0.05)^{50}} = \$765,283.49$$

Alternatively, we can recognize these cash flows as a growing annuity, whose present value can be found using equation 2.10.

$$Value_0 = \frac{30,000}{0.05-0.02} \times \left(1 - \left(\frac{1+0.05}{1+0.02}\right)^{-50}\right) = \$765,283.49$$

The present value of our lifetime earnings is $765,283.49. We should be indifferent between (i) receiving the after tax cash flows from working, and (ii) a check today in the amount of $765,283.49. If we receive the latter, we could invest it at our opportunity cost, 5% per year, and withdraw each year the same amount of money as what we would receive in after-tax earnings.

2.6.2 Go to College and Pay with our Savings

Figure 2.19 presents the cash flows if we go to school.

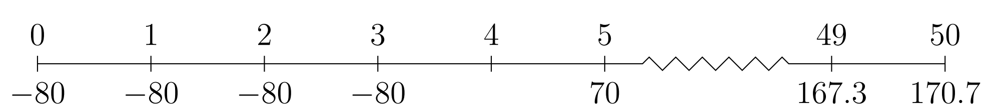

Figure 2.19: Go to College Timeline

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

$$Value_0 = -80,000 + \frac{-80,000}{1+0.05} + \frac{-80,000}{(1+0.05)^2} + ... + \frac{170,650}{(1+0.05)^{50}} = \$1,115,815.13$$

Copyright © 2020 Michael R. Roberts

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

$$80,000 + \frac{80,000}{0.05} \times \left(1 - (1+0.05)^{-3}\right) = \$297,859.84.$$

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 2.10) yields

$$Value_4 = \frac{70,000}{0.05 - 0.02} \times \left(1 - \left(\frac{1+0.05}{1+0.02}\right)^{-(50-4)}\right) = \$1,718,330.77.$$

Note, that we have only 46 years of earning, instead of 50, because we went to school for four years. Because $1,718,330 is the value of our after-tax earnings as of four years from today, we need to discount this amount back to today.

$$Value_0 = \frac{Value_4}{(1+0.05)^4} = \frac{1,718,330.77}{(1+0.05)^4} = \$1,413,674.97$$

The difference between the present values of the school costs and our earnings is 1,413,674.97-297,859.84 = $1,115,815.13, the same as we found above.

To summarize,

- If we don't go to college, the present value of our after-tax earnings is $765,283.49.
- If we go to college and pay for it with savings, the present value of our after-tax earnings is $1,115,815.13.

Going to school is the best decision from a financial perspective; it generates the most value (1,413,674.97-1,063,143.33=$350,531.64). This will not always be true. There are tradeoffs to each decision.

2.6.3 The Costs and Benefits of Going to School

Figure 2.20 shows the costs and benefits involved in choosing between skipping school and going to college and paying for it out of savings. The first four years shows the costs of going to college. We have to pay for college, and we miss out on work earnings. Only five years from now do we begin to see the payoff from college in the form of higher earnings.

Copyright © 2020 Michael R. Roberts

Years	0	1	2	3	4	5	50
Go to School	−80.0	−80.0	−80.0	−80.0	0.0	70.0	170.6
Skip School	0	30.0	30.6	31.2	31.8	32.5	79.1
Difference	−80.0	−110.0	−110.6	−111.2	−31.8	37.5	91.5

Figure 2.20: Cash flows for Going to School, Skipping School, and the Difference ($000s)

The present value of the costs of going to school is

$$80,000 + \frac{110,000}{(1+0.05)} + \frac{110,600}{(1+0.05)^2} + \frac{111,212}{(1+0.05)^3} + \frac{31,836}{(1+0.05)^4} = \$407,340.23.$$

These cash flows correspond to the tuition and fees for school *and* the missed earnings from working. The present value of the benefits of going to school is

$$\underbrace{\left[\frac{37,527}{(1+0.05)} + \frac{38,277}{(1+0.05)^2} + ... + \frac{91,485.44}{(1+0.05)^{46}}\right]}_{Value_4} \times \frac{1}{(1+0.05)^4} = \$757,871.86.$$

These cash flows correspond to the increased future earnings. The difference of these two present values gets us the net present value of going to school relative to skipping school: 757,871.86 - 407,340.23 = $350,531.64. Note, this is the same number we computed above. Intuitively, the decision of whether or not to go to school is a question of whether the increase in our future earnings will be greater than the short-term costs of school *and* foregone work earnings.

2.7 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 with time periods above the line - "0" for today, "1" for next period, "2" for the following period, and so on. Be sure to note the time period units. In this chapter, we only examined periods of one year. Soon we'll explore other frequencies such as quarter, month, and day.

2. Below the line, list the cash flows, *all of them*. In the applications above, we 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 helps avoid confusion in more complex settings.

Copyright © 2020 Michael R. Roberts

3. Identify the discount rate we'll use to discount or compound the cash flows (i.e., to change the time unit of the cash flows so we can add or subtract cash flows at different points in time.)

For those familiar with spreadsheet programs (e.g., Excel, Google Sheets), timelines should look familiar. They're nothing more than handwritten spreadsheets. So, timelines serve 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 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.

2.8 Key Ideas

We saw in Chapter 1 that you need to know one equation to do most everything in finance, our fundamental value relation.

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

Every equation we used in this chapter is just a special case of this one. Annuities (equation 2.5) and perpetuities (equation 2.8) are the cases where the cash flows are all the same; $CashFlow_1 = CashFlow_2 = CashFlow_3 = ...$ Growing annuities (equation 2.10) are the cases where the cash flows are all growing at a constant rate; $CashFlow_2 = CashFlow_1 \times (1+g), CashFlow_3 = CashFlow_1 \times (1+g)^2, ...$ Annuity cash flows (equations 2.6 and 2.12) are just algebraic rearrangements of these equations.

So, we've just been using our fundamental valuation relation all along. Importantly, our fundamental valuation 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.

Copyright © 2020 Michael R. Roberts

- Money has a time unit determined by when it is received or paid. If we want to add or subtract money, we have to make sure it has the same time unit and the discount factor, $1/(1+r)^t$, is our exchange rate for time. Multiplying money by the discount factor (i.e., dividing by $(1+r)^t$ moves money t-periods *back* in time. Dividing by the discount factor (i.e., multiplying by $(1+r)^t$) moves money t-periods forward in time.

- Net present value (NPV) is the present value of all benefits minus the present value of all costs. 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. (Yes, we discussed this in the previous chapter but it's such a central concept for financial decision making that it's hopefully worth the repetition.)

- 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 (or non-growing when $g = 0$) annuity at time t is

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

  where CF is the first cash flow one period from today, 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). Note the special case when $r = g$. We can't use the top result 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. These two effects cancel and we're left with $CF/(1+r) \times (T-t)$.

Copyright © 2020 Michael R. Roberts

- The present value of a growing (or non-growing when $g = 0$) perpetuity at time t is:

$$Value_t = \frac{CF}{r-g}$$

where CF is the first cash flow one period from today, 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 less than the discount rate. If not, the present value is infinite.

2.9 Technical Appendix

This appendix derives the shortcut formulas for common cash flow streams, specifically, annuity (equation 2.5) and perpetuity (2.8) and their constant growth versions (equations 2.10 and 2.8). Let's focus on a growing annuity because each result is just a special case.

$$\begin{aligned}&\frac{CF}{1+r} + \frac{CF(1+g)}{(1+r)^2} + ... + \frac{CF(1+g)^{T-1}}{(1+r)^T} \\ &= \frac{CF}{1+g} \underbrace{\left(\frac{1+g}{1+r} + \frac{(1+g)^2}{(1+r)^2} + ... + \frac{(1+g)^T}{(1+r)^T} \right)}_{\text{Partial sum of geometric series}}\end{aligned} \tag{2.13}$$

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

$$\sum_{t=1}^{T} a_t = a_1 \left(\frac{1+c^T}{1-c} \right) \tag{2.14}$$

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 2.14 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 2.14 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

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

Copyright © 2020 Michael R. Roberts

Plugging this result into the parenthetical term in equation 2.13 produces the present value of a growing annuity result, equation 2.10.

$$\frac{CF}{r-g}\left(1-\left(\frac{1+g}{1+r}\right)^{T}\right)=\frac{CF}{r-g}\left(1-\left(\frac{1+r}{1+g}\right)^{-T}\right) \tag{2.15}$$

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 2.14 does not hold. However, looking at equation 2.13, 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 2.15 reduces to our annuity result, equation 2.5

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

- When the number of time periods T gets arbitrarily large we have our growing perpetuity result, equation 2.8

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

  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 2.8

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

2.10 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.

Copyright © 2020 Michael R. Roberts

a. 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?

b. 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 (*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?

2.5 (*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?

2.6 (*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?

2.7 (*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")?

Copyright © 2020 Michael R. Roberts

2.8 (*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%?

2.9 (*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.

2.10 (*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?

Copyright © 2020 Michael R. Roberts

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?

2.11 (*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?

2.12 (*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?

2.13 (*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.)

Copyright © 2020 Michael R. Roberts

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

2.14 (*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.

2.15 (*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?

2.16 (*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?

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

(a) \$70,000 every year starting next year for 20 years.

Copyright © 2020 Michael R. Roberts

(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?

2.18 (*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?

2.19 (*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%.

2.20 (*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?

2.21 (*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?

Copyright © 2020 Michael R. Roberts

2.22 (*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?

2.23 (*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?

2.24 (*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?

Copyright © 2020 Michael R. Roberts

2.25 (*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. If their annual salary grows at 7%, what is the present value of their compensation?

b. What is the future value of their salary at the end of 10 years?

c. What is the present value of the banker's compensation, if they also receive an annual bonus equal to 125% of their annual compensation?

2.26 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?

Copyright © 2020 Michael R. Roberts

2.27 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.

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%.

Copyright © 2020 Michael R. Roberts

	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 2.21: Two-Way Data Table

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

2.28 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 Unition.

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?

Copyright © 2020 Michael R. Roberts

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?

2.29 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.

Copyright © 2020 Michael R. Roberts

Chapter 3

Financing Purchases

Fundamental value relation

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

Sometimes we don't have enough money to buy something so we have to **finance** the purchase. That is, we have to get money from someone else. But, how do we know when to turn to someone to get money? And, what are the implications of getting money from someone else?

This chapter

- 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 yield curve, and
- applies our fundamental value relation to answer several questions including:
  - How can we assess a college payment plan?
  - How does an auto lease work and can we recognize it as a loan?

Copyright © 2020 Michael R. Roberts

- How should we use credit cards?
- Are lower loan interest rates always better than higher loan interest rates?
- How does the federal reserve bank (i.e., "The Fed") impact asset values and the broader economy through its monetary policy?

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 be careful with how we measure the discount rate, r. We also have to recognize that r can vary depending on when the cash flow arrives. That is, cash flows at different points in time can have different discount rates. Otherwise, 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 Compounding

Chapter 2 focused on examples in which the cash flows arrived and departed at an annual frequency. We also assumed interest was paid once a year. The real world is littered with examples in which these assumptions are untrue.

Imagine that Bank of America offers a savings account with an Annual Percentage Rate (APR) 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. The \$110 is the future value of \$100 one year from today at a 10% annual interest rate. The computation is illustrated in figure 3.1.

Figure 3.1: Annual Compounding

What if the bank pays us interest twice a year, once every six months? This scenario and the corresponding interest computations are illustrated in figure 3.2.

A period is now six months. The **periodic interest rate** measures how much interest is paid each period and calculated as the APR divided by the number of periods or $0.10/2 = 0.05$. After one year we have \$110.25, \$0.25 more than what we had with annual compounding. This slight 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

Copyright © 2020 Michael R. Roberts

6-months	0	1	2
Cash flows	100	\$105	\$110.25
		$= 100\left(1+\frac{0.10}{2}\right)$	$= 105\left(1+\frac{0.10}{2}\right)$

Figure 3.2: Semi-annual Compounding - Periodic Rate

in the second half of the year, which is $5 \times 0.05 = \$0.25$. This process of interest earning interest is often referred to as **compound interest**.

Another way to tackle this problem is to find an **effective annual rate** or **EAR**. The EAR measures how much interest is earned in a year. Because of the semi-annual compounding, the EAR will differ from the APR of 10% in this example. Figure 3.3 illustrates the problem using the EAR.

6-months	0	1	2
Cash flows	100		\$110.25
			$= 100\left(1+\frac{0.10}{2}\right)^2$

Figure 3.3: Semi-annual Compounding - Effective Annual Rate

We multiply our *original deposit*, \$100, by $(1+0.10/2)^2$ to determine how much money we have in the bank at the end of the year. Over the year, our money earned $(1+0.10/2)^2-1 = 0.1025$. Thus, our EAR is 10.25%, slightly larger than the APR because of the semi-annual (twice per year) compounding.

To hammer home the point, and highlight a pattern, let's see what happens with monthly compounding, which is illustrated in Figure 3.4. Our monthly interest rate is $0.10/12 = 0.0083$, which measures how much interest is earned each month. Our effective annual interest rate is $(1+0.10/12)^{12} - 1 = 0.1047$. This rate measures how much we earn over the entire year including the effects of compounding. Regardless of whether we use the monthly interest rate or effective annual interest rate, we get the same ending balance of \$110.47. This value is higher than the balance we would have if interest is compounded annually or semi-annually. Why? Because more interest is earning interest.

3.1.1 Summary

Let's summarize and generalize what we've done.

Copyright © 2020 Michael R. Roberts

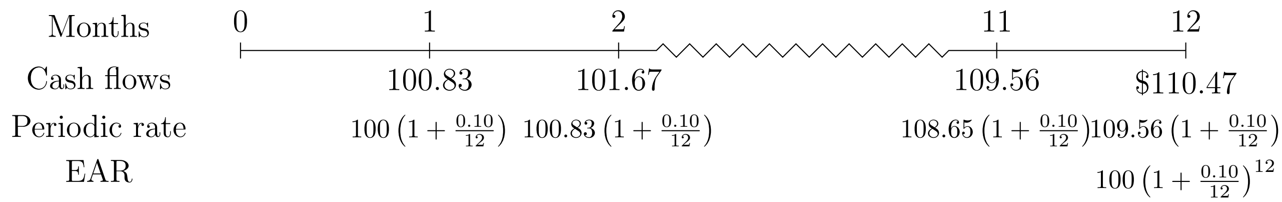

Figure 3.4: Monthly Compounding - Periodic Interest Rate

1. **Compounding frequency** refers to the number of times per year interest is compounded (i.e., paid). Here is a table of some common compounding frequencies and the corresponding number of times per year interest is compounded, which I denote with the letter k. Admittedly, decennial and bi-annual compounding are uncommon,

Compounding Frequency	Times per Year (k)
Decenial	0.10
Bi-annual	0.5
Annual	1
Semi-annual	2
Quarterly	4
Monthly	12
Trading days	252
Daily	365

Table 1: Common Interest Compounding Frequencies

   but they are included to emphasize the compounding frequency is measured as times per year.

2. Three types of interest rates include the following.

   (a) The **annual percentage rate** (**APR**) measures the **simple interest**, or interest ignoring compounding, earned in a year. 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** (i) measures the interest earned in a compounding

[1]The 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.

Copyright © 2020 Michael R. Roberts

period (e.g., quarter, month, day).

$$i = \frac{APR}{k}. \tag{3.1}$$

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

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

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

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,

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

Let's look at some applications.

3.2 Application: Paying for College

3.2.1 The Penn Payment Plan

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 $320,000 \div 48 = \$6,666.67$ per month beginning at the start of school. The timeline of tuition costs under the plan is illustrated in figure 3.5.

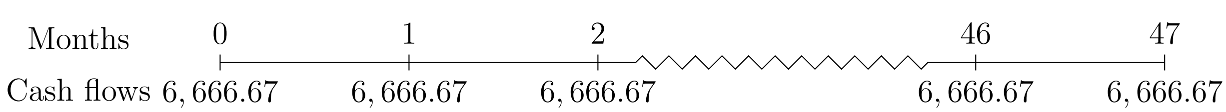

Figure 3.5: College Tuition Payment Plan

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 value of the tuition payments.

Copyright © 2020 Michael R. Roberts

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

$$i = \frac{APR}{k} = \frac{0.05}{12} = 0.004167.$$

We can discount the tuition costs with this rate *and* by measuring time not in years but in periods, months in this case.

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

Notice a few things. The first payment is made today and is therefore not discounted. The exponents in the denominators range from 1 to 47 and correspond to months.

While this is easily computed in a spreadsheet, we should recognize that the second through 48th payments are a monthly annuity. As such, we can use our annuity result (equation 2.5) to compute the present value of these cash flows where we measure time in periods and use the periodic interest rate to discount them.

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

Alternatively, we can compute this present value using the EAR and measuring time in years. The EAR (r) is

$$r = (1+i)^k - 1 = (1+0.004167)^{12} - 1 = 0.051162$$

The present value calculation looks as follows.

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

However, we *cannot* use our annuity formula with the EAR and measuring time in years.

This amount is less than the present value of the annual \$80,000 payments (\$297,859.84) we computed in the Chapter 2 for two reasons. First, the payments 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 as of 2022

Copyright © 2020 Michael R. Roberts

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 $1,200 \div 24 = \$50$.

Be careful with zero 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]

3.2.2 Financing College

In chapter 2, we asked whether it made financial sense to go to college, assuming we had the money our education expenses. Let's revisit that question assuming we must borrow the money to pay for school. Recall our assumptions.

- School costs $80,000 per year payable at the start of each year for four years.
- If we skip school, we will work for 50 years earning $30,000 after tax at the end of the first year. This amount will grow by 2% per year thereafter.
- If we go to school, we will work for 46 years earning $70,000 after tax one year after graduating. This amount will grow by 2% per year thereafter.
- Our opportunity cost of capital is 5% per year.

Let's assume that the loan APR is 7% with annual compounding and payments, and we don't start repaying the loan until five years from now. Annual compounding implies that the APR, periodic interest rate, and EAR are all equal.[3] Let's also assume that we have 10 years to repay the loan starting after we graduate. Finally, we'll assume that interest accrues on the loan while we're in school, meaning the loan balance will grow by the loan interest rate while we're in school.

How much do we need to borrow? Let's keep it simple and assume we need $4 \times 80,000 = \$320,000$ today, which the school will hold for us and use to make the annual payments. Having the school hold the money means we can't invest this borrowed money, which is

[2]For 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]Interest on most student loans is compounded monthly to coincide with monthly payments. We're abstracting from that complication to focus on the important issues.

Copyright © 2020 Michael R. Roberts

typical of most student loans in which the lender makes payments directly to the school. At the end of school four years from today, we'll owe our lender

$$320,000 \times (1 + 0.07)^4 = \$419,454.72$$

because interest accrues while we're in school.

The annual loan payments we'll need to make once we graduate are illustrated in figure 3.6. The CF are the constant annual payments we need to make on the loan. As such, the

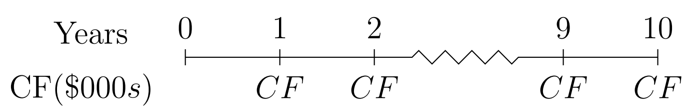

Figure 3.6: Loan Payments

sum of their present values equals the loan principal - how much we are borrowing.

$$419,454.72 = \frac{CF}{(1+0.07)} + \frac{CF}{(1+0.07)^2} + ... + \frac{CF}{(1+0.07)^{1}0} \tag{3.4}$$

Because the loan payments form an annuity, we can rewrite equation 3.4 as follows.

$$419,454.72 = \frac{CF}{0.07}\left(1 - (0.07)^{-10}\right)$$

Solving for CF produces

$$CF = \frac{419,454.72 \times 0.07}{1 - (1+0.07)^{-10}} = \$59,720.92.$$

(This is just our annuity cash flow formula, equation 2.6, from chapter 2.) So, one year after we graduate - five years from today - we'll receive income from our job, \$70,000, and have to pay \$59,720.92 to the bank. While our income continues to grow at 2% per year, we'll continue to make loan payments for another nine years, after which the only cash flows are our after-tax income.

Figure 3.7 presents the timeline for our problem. All of the cash flows are in \$000s. There are no cash flows associated with school because the lender pays tuition and fees directly to the school; nothing comes in or out of our pocket. Starting in year 5, we earn income (\$70,000) from which we must make our loan payment of \$59,720.92. We continue paying off the loan for nine more years - 10 in total - after which we no longer have any more loan payments.

Copyright © 2020 Michael R. Roberts

Years	0	1	4	5	6	14	15	49	50
Earnings				70.0	71.4	83.7	85.3	167.3	170.6
Loan Pmts.				59.7	59.7	59.7			
Difference	0	0	0	10.3	11.7	24.0	85.3	167.3	170.6

Figure 3.7: Go to College with Financing Timeline

To get the value of going to school when borrowing money, we can discount and sum the earnings and loan payments and then take the difference.

$$\begin{aligned}\text{Present value of earnings} &= \frac{70,000}{(1+0.05)^5} + \frac{71,400}{(1+0.05)^6} + ... + \frac{170,650}{(1+0.05)^{50}} \\ &= \$1,413,674.97\end{aligned}$$

$$\begin{aligned}\text{Present value of payments} &= \frac{59,721}{(1+0.05)^5} + \frac{59,721}{(1+0.05)^6} + ... + \frac{59,721}{(1+0.05)^{14}} \\ &= \$379,388.49\end{aligned}$$

The difference is 1,4113,674.97 - 379,388.49 = \$1,034,286.48. Alternatively, we can discount and sum the differences between the earnings and loan payments directly.

$$\begin{aligned}\text{Present value of difference} &= \frac{10,279}{(1+0.05)^5} + ... + \frac{23,936}{(1+0.05)^{14}} + ... + \frac{170,650}{(1+0.05)^{50}} \\ &= \$1,034,286.48\end{aligned}$$

Using the Annuity Formulas

Because of their prevalence in practice, it's useful to show how to solve this same problem using the annuity formulas (equations 2.5 and 2.10). The cost of the loan today (i.e., present value) is:

$$Value_0 = \underbrace{\frac{59,720.92}{0.05} \times \left(1 - (1+0.05)^{-10}\right)}_{Value_4} \times \frac{1}{(1+0.05)^4} = \$379,388.49$$

Notice two things about this calculation. First, we compute the present value of the annuity as of period 4, which we discount back to period 0 by dividing by $(1+0.05)^4$. Second, we compute the cost of the loan to us using *our* opportunity cost of capital, 5%, not the loan rate, 7%. Our opportunity cost is what matters for valuing the loan because if we didn't have to repay the loan and instead could invest the money, we would only earn 5%.

Copyright © 2020 Michael R. Roberts

The present value of earnings is computed similarly.

$$Value_0 = \underbrace{\frac{70,000}{(0.05-0.02)} \times \left(1 - \left(\frac{1+0.05}{1+0.02}\right)^{-46}\right)}_{Value_4} \times \frac{1}{(1+0.05)^4} = \$1,413,674.97$$

The difference between the present values of our earnings and the cost of the loan is 1,413,674.97 - 379,388.49 = $1,034,286.48, the same as above.

To summarize,

- If we go to college and borrow money to pay for it, the present value of our after-tax earnings less the cost of college and the loan is $1,034,286.48.
- If we don't go to college, the present value of our earnings is

$$\frac{30,000}{(1+0.05)} + \frac{30,600}{(1+0.05)^2} + ... + \frac{79,164.35}{(1+0.05)^{50}} = \$765,283.49$$

Going to school, even if we have to borrow, is financially advantageous by over a quarter million dollars (1,034,286.48 - 765,283.49 = $269,002.99).

However, look at figure 3.7. Five years from today, after we've graduated and begun working, we'll only have $10,279.08 on which to live? It's not clear that will cover all of our necessary (i.e., **nondiscretionary**) expenses - things like rent, food, clothing, etc. So, while going to school may be very valuable, an additional benefit of our analysis is in highlighting whether it is *feasible*. In this example, we may have to speak with our lender about deferring payment for a few years until our salary is sufficiently large. Or, we may have to find an alternative school major and career to ensure we can afford to repay our loan.

3.2.3 The Costs and Benefits of Going to School by Borrowing Money

Figure 3.8 presents another perspective on the school choice problem. The first row presents the cash flows of skipping school, i.e., our earnings. The second row presents the cash flows of going to school by borrowing money (the difference row in figure 3.6). The third row presents the difference between these two sets of cash flows and correspond to the incremental cash flows of going to school with borrowed money relative to not going to school.

The first four years show the costs of going to school - foregone earnings. Years five through 14 show the cost of borrowing; we have a loan to repay. The present value of these

Copyright © 2020 Michael R. Roberts

Years	0	1	2	3	4	5	14	15	50
Skip	0	30	30.6	31.2	31.8	32.5	38.8	39.6	79.2
Go & Borrow	0	0	0.0	0.0	0.0	10.3	23.9	85.3	170.6
Difference	0	30	30.6	31.2	31.8	22.2	14.9	−45.7	−91.5

Figure 3.8: Cash flows for Skipping School, Financing School, and the Difference ($000s)

costs is

$$Value_0 = \frac{30,000}{(1+0.05)} + \frac{30,600}{(1+0.05)^2} + ... + \frac{14,900}{(1+0.05)^{14}} = \$229,897.98.$$

The benefits of higher earnings are experienced in years 15 through 50. The present value of these benefits is

$$Value_0 = \frac{45,745}{(1+0.05)^{15}} + \frac{46,660}{(1+0.05)^{16}} + ... + \frac{91,485}{(1+0.05)^{50}} = \$498,900.97.$$

The difference between the present values of the costs and benefits is the NPV of going to school and financing it, 498,900.97-229,897.98=$269,002.99, exactly what we saw above.

3.3 Application: Auto Lease

A lease is a loan permitting the lessee to use the lessor's asset, such as a car or airplane. 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 start an airline and we need a plane. If we don't have enough money, we borrow money from a bank. But, maybe we don't want to own the plane, we just want to use it for a while. In this case, we can borrow an airplane from our friend that owns one in exchange for periodic payments. This is a lease. In this scenario, we are the lessee who is borrowing the plane, our friend is the lessor who owns the plane. Leases range in duration from several months to several years. At the end of the lease, we return the asset or, in some cases, we may have the option to purchase the asset. The key difference between a lease and a loan is that in a lease we don't actually own the asset. We're just borrowing it. Let's look at an example.

As of 2021, a new 2021 McLaren 720s Spider came with the following lease terms.

Copyright © 2020 Michael R. Roberts

- The length or term of the lease is three years.
- A \$35,000 payment must be made at the closing of the deal (start of the lease). This payment is sometimes referred to as "cash out of pocket."
- Monthly payments of \$4,820 - not including taxes - must be made over the three-year lease term.

Let's figure out the interest rate implied by these terms.

To do so, we'll need to understand how much we're borrowing, which requires two pieces of information: (i) the price of the car and (ii) the residual value of the car. A nicely appointed version of this car retails for about \$360,000. (I can dream.) The residual value is an estimate of what the car will be worth when we return it in three years. The difference between the cost and the residual value is how much we are borrowing.

Assume the residual value is 60% of the purchase price, or $0.60 \times 360,000 = \$216,000$. (The dealer should tell us this number.) In other words, when the lease ends in 2024, the car we return to the dealer will be worth \$216,000. Therefore, we need to borrow 360,000 − 216,000 = \$144,000 for three years.[4]

Figure 3.9 presents the timeline for the lease cash flows. When we start the lease, we are effectively receiving a \$144,000 loan from the dealer to take possession of the car. In return, we promise to pay \$35,000 to the dealer immediately. This cash out of pocket reduces the size of the loan from \$144,000 to \$109,000. We also promise to make 36 monthly payments of \$4,820 and return the vehicle three years from today. (Ignore the option to purchase the vehicle at the end of the lease.)

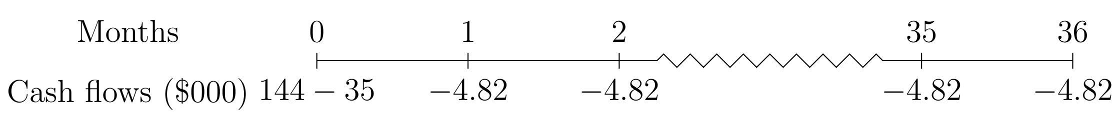

Figure 3.9: Auto Lease Payments

We can find the implied interest rate on the loan using our fundamental value relation

[4] For those paying close attention, subtracting the \$216,000 from the \$360,000 is a no-no because they have different time units. In this case, this error works in our (the lessee's) favor because it reduces the size of the loan and the corresponding monthly payments.

Copyright © 2020 Michael R. Roberts

with the periodic (monthly) discount rate, i.[5]

$$144,000 - 35,000 = 109,000 = \frac{4,820}{(1+i)} + \frac{4,820}{(1+i)^2} + ... + \frac{4,820}{(1+i)^{35}} + \frac{4,820}{(1+i)^{36}}$$

Because the monthly lease payments are an annuity, we can use equation (2.5) to express the present value.

$$109,000 = \frac{4,820}{i} \times \left(1 - (1+i)^{-36}\right)$$

Solving for i with a spreadsheet program yields $i = 0.0277$. In other words, we're paying 2.77% *per month* in interest. The APR on the lease is $12 \times 0.0277 = 0.3319$, or 33.19%. As of 2022 in the U.S., that's an extremely high interest rate to pay on a loan.

Now, maybe we've underestimated the price of the car or overestimated its residual value. Either mistake would underestimate how much we're borrowing and consequently inflate the interest rate. (Try changing the numbers to see this.) We could have also mistakenly assumed that the $35,000 up front payment is all down payment, as opposed to covering fees and taxes. Regardless, these numbers raise some questions to be asked when we're sitting down with the dealer.[6]

3.4 Application: Financing a Home with a Mortgage

A mortgage is nothing more than a loan used to purchase property. Home loan, real estate loan, and mortgage are all the same thing. Consider an example.

We found a home that costs $625,000. To pay for it, we'll put down 20% which means we'll write a check for 625,000×0.20 = $125,000, and borrow the rest, $500,000.[7] Now we have to choose a mortgage product, and there are lots. The most popular mortgage in the U.S. is a fixed-rate, 30-year mortgage, which requires borrowers pay a fixed amount every month for 30 years.[8] The mortgage APR is 3% compounded monthly to match the payment frequency.

[5]You should always ask for the lease APR to understand how much interest you're paying. Often, dealers will mention a **money factor** instead of an APR. Fortunately, it's easy to compute the APR by multiplying the money factor by 2400.

[6]There are often lots of fees and taxes buried in these financing arrangements. A recent car lease my wife entered into revealed a tax rate of 9%. There were bank fees, taxes on these bank fees, a tax on the down payment, registration fees, encumberance fees, electronic filing fees, and taxes on the monthly payments! This adds up to a lot of additional expenses that are unique to leasing the car.

[7]Paying 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 mortgage insurance, which protects the lender in case the borrower defaults on, that is doesn't repay, the loan.

[8]Alternative mortgage products include different maturities - 10, 15, and 20-year - and adjustable rate mortgages (ARMs) whose interest rate, and therefore monthly payment, changes every year.

Copyright © 2020 Michael R. Roberts

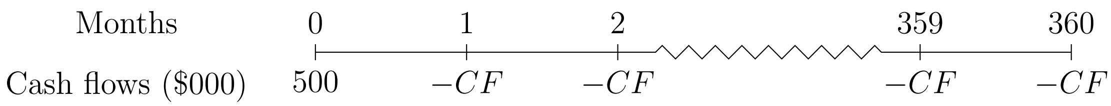

Figure 3.10: Mortgage

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.

Because the cash flows are the same, equally spaced out in time, and for a finite amount of time, they correspond to an annuity. So, we can use equation 2.6 to compute the annuity cash flow, which in this case is just the monthly mortgage payment. But, we have to measure time in periods, or months in this case, 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.

3.4.1 Amortization Table

Accompanying most mortgages is an amortization table containing monthly mortgage information for the life of the loan. 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.

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

Copyright © 2020 Michael R. Roberts

Each period is one month in duration. 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.

During each month, interest accumulates 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 500,000×0.0025 = $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 balance

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. In other words, apply the fundamental valuation relation to the remaining mortgage payments using the loan interest rate. Because the loan payments make up an annuity, we can use equation 2.5. For example, after 5 years, or the 60th payment, we would owe the bank

$$Principal_{60} = \frac{2,108.02}{0.0025} \times \left(1 - (1 + 0.0025)^{-(360-60)}\right) = \$444,531.82. \tag{3.5}$$

The number of periods, 360 − 60 = 300, is the number of remaining payments.

Copyright © 2020 Michael R. Roberts

Payment Composition

The amortization table shows that each monthly mortgage payment is comprised of two pieces: interest and principal reduction. While the payments are constant, the proportions of each payment used to pay interest and reduce principal change every month. We can see this in the amortization table above, table 2.

Consider the 100th mortgage payment. The loan principal *before* the payment is made equals

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

The interest accruing in the 100th month is 403,756.28×0.0025=\$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÷2,108.02 = 47.9% of our 100th mortgage payment is going towards paying interest, 1,098.63÷2,108.02 = 52.1% is going towards principal reduction.

3.4.2 Refinancing a Mortgage

One of the biggest mistakes homeowners make is not refinancing their mortgage when interest rates decline. To illustrate the value gain that can be achieved from refinancing, consider the following example. Imagine that five years have elapsed so "today," period 0, is just after we've made our 60th mortgage payment. Let's also assume that the interest rate on a mortgage "refi" is 2.75%, 25 basis points lower than the 3.00% we are paying on our mortgage. (One **basis point** is 0.0001 or 0.01%; one **point** is 0.01 or 1%.) Let's see how much money we can save by refinancing our mortgage to take advantage of this lower rate.

Step 1: Compute how much is owed to the bank. We did this above (equation (3.5)) by discounting the remaining cash flows by the loan interest rate and summing. After five years, we owe the bank \$444,531.82.

Step 2: Revalue the loan at the new interest rate. The new interest rate of 2.75% implies a new periodic interest rate of 0.0275÷12 = 0.002291. With 300 payments remaining, the current *value* of our mortgage at the *new* interest rate is

$$Value_0 = \frac{2,108.02}{0.0275/12} \times \left(1 - (1 + 0.0275/12)^{-(360-60)}\right) = \$456,963.05.$$

Note that by using the current market interest (2.75%) rate to discount cash flows we get the *value* of our loan. Discounting by the loan interest rate (3.00%), we get the *principal*

Copyright © 2020 Michael R. Roberts

outstanding on our loan - what we owe. Because interest rates have changed, the two are not equal.

Step 3: Subtract what is owed to the bank from the new value of the mortgage. The difference between the value of the loan and the amount we owe on the loan is 456,963.05−444,531.82=$12,431.23. This is the present value, as of five years into the loan, of refinancing our mortgage. Because its positive, it is beneficial to refinance our loan.

Step 4: Compare refinancing savings with refinancing costs. There are often fees associated with refinancing a mortgage - appraisal, doc, loan, etc - that are charged at the time of refinancing. As long as these costs are less than the benefits of refinancing, then refinancing makes sense. For example, if the total cost of refinancing our mortgage is $2,500, then the net present value (NPV) of the refinancing is

$$NPV = 12,431.23 - 2,500 = \$9,931.23.$$

The positive NPV implies that refinancing our mortgage will create $9,931.23 of value for us today. We should refinance.[9]

Intuition

The procedure above hides some intuition. For example, why does the difference between the market value of the loan and the loan principal equal the present value of the interest savings? Let's walk through what's going on with the cash flows when we refinance. The timeline for the original mortgage after five years has passed is shown in figure 3.11. Note that there are only 300 months - 25 years - of payments remaining.

Months	0	1	2	...	299	300
Cash flows ($000)		−2.108	−2.108		−2.108	−2.108

Figure 3.11: Existing Mortgage - Remaining Payments after Five Years

If we refinance into a new 30-year mortgage, our new monthly payments will be

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

The new timeline is shown in 3.12.

There are two differences between the old and new mortgages:

[9] Because refinancing is sometimes costly, we also need to consider the value of waiting to refinance in case the interest rate falls further. However, if you're like me and most everyone else, you can't predict what interest rates are going to do so consideration of this option is a luxury most of us can ignore.

Copyright © 2020 Michael R. Roberts

Months	0	1	2	...	359	360
Cash flows ($000)		−1.814	−1.814		−1.814	−1.814

Figure 3.12: New Mortgage Payments

1. The new mortgage has lower monthly payments because (i) less is owed to the bank ($444,531.82 vs. $500,000), and (ii) the interest rate is lower (2.75% vs. 3.00%).

2. The new mortgage has more monthly payments - 360 versus 300. (This need not be the case as we discuss below.)

Figure 3.13 overlays the two timelines and subtracts the old cash flows from the new ones. All cash flows are in thousands of dollars.

	0	1	300	301	359	360
New Mort.	$444,5$	-1.814	-1.814	-1.814	-1.814	-1.814
Old Mort.	$-444,5$	$-(-2.108)$	$-(-2.108)$	-0	-0	-0
Diff.	$=0$	$=0.294$	$=0.294$	$=-1.814$	$=-1.814$	$=-1.814$

Figure 3.13: Costs and benefits of Refinancing

When we refinance the mortgage, we receive $444,531.82 from our new lender that is immediately transferred to our old lender to pay off the old mortgage. One month later, we save $294 because of the lower monthly payment on the new mortgage. These savings continue for 300 months until the old mortgage would have been fully paid off. However, the new mortgage has an additional 60 months of payments. So, the benefits of the new mortgage are 300 months of savings; the costs are an additional 60 months of payments.[10]

Now let's compute the present values of the costs and benefits by discounting the corresponding cash flows by the new interest rate. We use the new interest rate because it is the current market rate, which is what's relevant for valuation. The old loan interest rate is just that - an old quoted interest rate that is useful for identifying the outstanding principal on the original loan but not much else. The benefits are a 300 month annuity of $293.26 in savings whose present value is

$$Value_0 = \frac{293.26}{0.0275/12} \times \left(1 - (1 + 0.0275/12)^{-300}\right) = \$63,570.63.$$

[10] When refinancing, we don't have to extend the length of the mortgage. We could have refinanced into a shorter maturity mortgage, say 10- or 20-years. But, this shorter maturity would come with higher monthly payments since we have less time to pay down the mortgage.

Copyright © 2020 Michael R. Roberts

The costs are a 60 month annuity of $1,814.76 starting 301 months from today and whose present value is

$$Value_0 = \underbrace{\frac{1,814.76}{0.0275/12} \times \left(1 - (1+0.0275/12)^{-(360-300)}\right)}_{Value_{300}} \times \frac{1}{(1+0.0275/12)^{300}} = \$51,139.40$$

Notice the annuity equation gives us the value of the costs as of 25 years from today (period 300). We must discount that value by $(1+0.0275/12)^{300}$ to get the present value of the costs.

The difference between the present values of the costs and benefits is the value of the refinancing decision, which equals $63,570.63 - 51,139.40 = \$12,431.24$, the same amount we computed above.

Some comments

- We could have chosen any maturity for our new mortgage, 10-, 15-,20-, or 30-year. As long as we stay in the house until it is fully paid off, the value of refinancing is the same.
- If we choose to sell our home any time before the 25 years remaining on our existing mortgage, then the value of refinancing is reduced. Intuitively, by selling the house early, and therefore repaying the mortgage early, we are not taking full advantage of the lower interest rate on the new loan. If we sell too quickly, the benefits of refinancing may be swamped by any costs of refinancing (e.g., fees). Therefore, refinancing only makes sense if we plan on staying in the home sufficiently long enough to realize the benefits from the lower mortgage payments.
- Decreases in interest rates lead to increases in the value of a mortgage. This follows immediately from our fundamental value relation. If r declines, and cash flows are unchanged, then value must increase and vice versa. This inverse relation between value and interest rates, or discount rates more generally, applies to *all* assets.

3.4.3 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 better than saving our money. We'll consider some popular arguments and then answer the question objectively with what we've learned thus far.

Copyright © 2020 Michael R. Roberts

To do so, we'll use the mortgage we've been analyzing thus far; the 30-year, fixed-rate loan for $500,000 with an APR of 3%. Let's also assume that we have $250 of extra money each month that we can either save or apply to our mortgage.

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, but over the same loan term. 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 2.9 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

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

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, just means that our last payment is a partial payment, which can be calculated as follows. Step one, estimate the present value of making 302 payments of $2,358.02.

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

This calculation shows that the present value of 302 payments is 500,000 - 499,472.39 = $527.62 short of paying off the loan in *today's dollars*. Step two, compute the future value of this shortfall 303 months in the future.

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

This is the last (partial) payment occurring in month 303 to fully pay off the loan.

Copyright © 2020 Michael R. Roberts

2. We'll pay less interest. This is true, but it is not a compelling argument because it ignores the opportunity cost of paying down the mortgage. Additionally, this argument is typically presented in a seemingly intuitive but completely nonsensical manner. The argument often goes like this.

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

$$302 \times 2,358.02 - 500,000 = \$212,122.04$$

in interest. Thus, we'll save 258,887.20 - 212,122.04 = $46,765.16 in interest.

Can you see what's wrong with these calculations? They violate rule #1 of finance: Don't add cash flows arriving or going 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.

Proper Analysis

The decision of whether or not to pay off our mortgage early is, at some level, simple from a financial standpoint.[11] If we can earn a higher rate of return on our savings than the interest rate we have to pay on our mortgage, then paying off our mortgage early is a bad idea and vice versa. In our mortgage example, as long as our money can earn more than 3%, directing extra money to paying off our mortgage more quickly is financially unwise.

[11] There are other financial considerations including the diversification of our wealth. Paying off our mortgage means we are accumulating equity in our home. Depending on how our other savings are invested, if we pay our mortgage down too quickly we can be overexposed to real estate risk. The 2008 financial crisis exposed just how risky real estate investing can be.

Copyright © 2020 Michael R. Roberts

Let's illustrate this by assuming that we can invest our money and earn a monthly compounded return of 5% per year on our savings. That is, our opportunity cost is 5%. Consider two strategies

1. Save our extra money

If we save the extra $250 per month for 30 years, we will have

$$Value_{360} = \underbrace{\frac{250}{0.05/12} \times \left(1 - (1 + 0.05/12)^{-360}\right)}_{Value_0} \times (1 + 0.05/12)^{360} = \$208,064.66$$

30 years from today. (Note: the annuity calculation computes the present value of our savings, which we compound to get the future value.)

2. Use extra money to pay down mortgage

Figure 3.14 shows what we can save if we redirect the $250 towards paying down the mortgage more quickly.

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

If we redirect extra money towards paying down the mortgage, we can't start saving until 303 months from today. At that time, we can apply all our money - the extra $250 *and* any money we would have had to pay towards our mortgage. In month 303, we make a partial mortgage payment equal to $1,124.31, meaning we have 2,108.02 - 1,124.31 = $983.71 that we can save plus the extra $250. 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

$$\begin{aligned} Value_{360} &= \underbrace{\left[1,233.71 + \frac{2,358.02}{0.05/12} \times \left(1 - (1 + 0.05)^{-57}\right)\right]}_{Value_{303}} \times (1 + 0.05/12)^{57} \\ &= \$152,919.95. \end{aligned}$$

To summarize, the value 30 years from today of our saving the extra money is $208,064.66. The value 30 years from today of paying down our mortgage with the extra money is

Copyright © 2020 Michael R. Roberts

$152,919.95. The net value 30 years from today of saving our money is therefore

$$208,064.66 - 152,919.95 = \$55,144.71.$$

The net present value (NPV) of saving our money is $55,144.71 \div (1+0.05/12)^{360} = \$12,342.85$. We should save our money.

Paying off a mortgage early is a tradeoff between saving now versus later. If our opportunity cost is greater than our mortgage cost, as it is in this example, then directing extra money towards savings creates value for us. If our opportunity cost is less than our mortgage cost, then directing extra money towards savings destroys value for us.

Table 1 showed that the average annual return on the stock market over the last 100 years has been 11%. So, the temptation is to think, "we can easily make more than 3% by investing in the stock market." This is a mistake! We have to pay off our mortgage. The stock market does not have to pay 11% per year. In other words, our mortgage payments are essentially risk-free, while investing in the stock market carries significant risk. There is risk mismatch between paying off a mortgage and investing in the stock market that makes this alternative an apples to oranges comparison.

What we need is a risk-free investment with a return greater than the cost of our mortgage. When we first buy our home and take out a mortgage, the cost of the mortgage will always be greater than the return we can earn on a risk-free investment because the bank doesn't see us as risk-free. We may default. As such, they charge us a higher interest rate because of that risk. This fact suggests that paying off a mortgage early is *always* a good idea. However, there are two exceptions.

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 cost of our mortgage. This is exactly what occurred in 2022. Rising inflation was met with large increases in interest rates on Treasury securities - the safest investment at the time. Many homeowners that had originated mortgages before the interest rate increases were able to take advantage of their low fixed rate mortgage by saving any extra money in Treasury securities and earning a higher return than the cost of their mortgage. (One person I know had a 2.75% 30-year fixed rate mortgage at the time the yield on 30-year Treasury bonds wsa over 4.00%.)

The other exception is if we, the homeowners, are willing to take some risk. For example, I asked 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, and how much money do we stand to lose if we do lose money?

Copyright © 2020 Michael R. Roberts

Using stock market data over the last 100 years, the probability that saving extra money turns out to be a bad decision is 26%. 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 very 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.

3.5 Application: Credit Cards

The previous two applications are perfect preparation for exploring how we should handle credit card debt. The basic lesson thus far has been: Use any money we can to pay off all debt that charges a higher interest rate than what we can earn saving that money in an investment of similar risk. Or, recognize any risk we may be taking by saving the money in an investment with greater risk to achieve a higher return.

3.5.1 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. According to Bankrate.com, as of March 2022, the current interest rate on a 30-year loan is about 4.6%; the typical interest rate on a credit card is 16.34%! As of March 2022 - or as of most periods in U.S. history - earning more than 16.34% per year requires taking large risks. Thus, paying off credit card, or any high interest, debt is often a smart financial strategy, where "high interest" simply means greater than what we can earn on other (near) risk-free investments.

Copyright © 2020 Michael R. Roberts

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, in addition to 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 is that opening new credit card accounts require credit checks that can negatively impact your credit score generated by the credit agencies - Experian, Equifax, and Transunion. Your credit score is an important determinant in whether you can get credit, such as home and auto loans, and what interest rate you will have to pay.

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.

Copyright © 2020 Michael R. Roberts

3.5.2 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 we owe, we'll actually wind up owing more and more each month! Look at the first few lines of the amortization table 3

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, we'll wind up paying down the credit card, just over a very long period of time. And, because the effective annual rate we're paying is $(1 + 0.1899/12)^{12} - 1 = 0.2073$, this strategy is almost surely a terrible idea because we're unlikely to find a risk-free investment offering that much return.

Copyright © 2020 Michael R. Roberts

	Start Balance	Interest	Payment	End 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 3: Credit Card Amortization Table

3.6 Term Structure of Interest Rates

One important way auto loans and mortgages differ is their length or term. Common auto loan terms are 3, 4, 5, and 6-years. Common mortgage terms are 10-, 15-, 20-, and 30-years. Each of these terms has its own, possibly different, interest rate. The relation between the term of a loan and the interest rate is referred to as the **term structure of interest rates**. The term structure of mortgage interest rates as of January 2022 is presented in table 4

Term	APR
10	2.500%
15	2.400%
20	2.875%
30	3.374%

Table 4: Term Structure of Mortgage Rates, January 2022 (Source: Bankrate.com)

The important feature of table 4 to note is that the interest rate varies depending on the term of the mortgage.

From where do these interest rates come? Supply and demand in the market for mortgages! Homeowners demand mortgages and financial institutions, like banks, supply them. More demand leads to higher interest rates; more supply leads to lower interest rates. That mortgage rates differ across terms suggests that equilibrium supply and demand are different for these different mortgage products. In other words, there appears to be greater demand relative to supply for 30-year mortgages than for 20-year mortgages, which is why the interest rate for the former (3.374%) is greater than that for the latter (2.875%).

Why would people prefer 30-year mortgages over 20-year mortgages? One reason is that they expect interest rates in the distant future to be higher than what they'll be in 20 years. By entering into a 30-year fixed-rate mortgage, we insure ourselves against high

Copyright © 2020 Michael R. Roberts

future interest rates that may arise after 20 years. Likewise, the interest rate on a 20-year mortgage is higher than that on a 10-year mortgage. If we're worried that interest rates may rise after 10-years, a 20-year mortgage ensures we don't have to worry about that increase because a fixed-rate mortgage **locks in** or fixes our interest rate today.

This logic leads to the **expectations hypothesis** of the term structure of interest rates. The expectations hypothesis says that the slope of the term structure is governed by expectations about future interest rates. When the term structure is upward sloping, i.e., long-term interest rates are higher than short-term interest rates, as is the case in table 4, then investors in mortgages expect future mortgage rates to increase. When the term structure is downward sloping, then investors in mortgages expect future mortgage rates to decrease.

Unfortunately, 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 these mortgages. Consequently, the term structure changes frequently and often in unpredictable ways.

3.6.1 Many Term Structures

Above we discussed the term structure for fixed-rate mortgages. However, there are many term structures, one corresponding to each type of loan in the economy. Table 5 presents the term structure for investment-grade (i.e., low-risk) corporate bonds, which are loans to corporations like Microsoft and Johnson & Johnson.

Term	Yield (APR)
2	0.86%
5	1.60%
10	2.54%
30	2.98%

Table 5: Term Structure of Corporate Bonds, November 2021 (Source: St. Louis Federal Reserve Economic Data)

We slipped in a new term at the top of the table, **yield**. We'll discuss yields in more detail in chapter 7. For now, it suffices to know that a yield is just the APR of the bond. Like mortgages, corporate bonds exhibit an upward sloping term structure suggesting that investors in corporate bonds expect future interest rates on these loans will increase.

Arguably the most important term structure is that for Treasury securities, which are loans to the U.S. federal government. The government spends a lot of money on things

Copyright © 2020 Michael R. Roberts

like social security (payments to retired workers), Medicare (health care benefits for senior citizens), defense, and infrastructure (e.g., bridges and roads). It raises money through taxes. When the taxes the government raises are less than the money it spends, the government has to borrow money, which it does by issuing Treasury securities to investors.

Table 6 presents the term structure for Treasury securities for January 6, 2022 and 1981 (for which some data is missing).

Term	January 2022 Yield (APR)	January 1981 Yield (APR)
1 month	0.05%	
2 month	0.06%	
3 month	0.08%	15.01%
6 month	0.22%	
1 year	0.40%	14.06%
2 year	0.78%	14.26%
3 year	1.04%	14.19%
5 year	1.37%	14.11%
7 year	1.55%	14.14%
10 year	1.63%	14.21%
20 year	2.05%	14.33%
30 year	2.01%	

Table 6: Term Structure of Treasury Securities, January 6 2021 and 1981 (U.S. Department of the Treasury)

Focusing first on 2022's interest rates, notice that they are low relative to other loans (mortgages and corporate bonds). Treasuries are the gold-standard for safe investments, so their interest rates are often viewed as the closest thing to a truly **risk-free** interest rate. Mortgages and corporate bonds, on the other hand, come with a non-negligible amount of **default risk**, risk that the borrower won't be able to repay the loan. The U.S. government is believed by investors to always pay its debts, which historically it has. As such, the interest rate on Treasury securities is typically lower than that on all other types of loans in the U.S.

More recent interest rates are also low relative to what they were in 1981. There are several reasons for these differences, most notably inflation was significantly higher back in 1981. The points worth emphasizing here are (i) interest rates can and do vary quite a bit, and (ii) when people talk about "the term structure," they are almost always talking about the Treasury term structure because the size of Treasury markets is so large, and Treasury securities are so important for the global economy.

Copyright © 2020 Michael R. Roberts

3.6.2 Implications

There are several implications of different interest rates for different term loans. One was discussed above, namely, the expectations hypothesis. The term structure may be able to give us insight into future interest rates. By extension, the term structure may also give us insight into future economic conditions. Interest rates tend to be lower when economic growth is slower (e.g., recessions) and higher when economic growth is faster (e.g., booms). Downward sloping yield curves in which future interest rates are expected to decline have been shown to portend economic slowdowns.

The term structure also provides different discount rates for cash flows that arrive or depart at different times. For example, if our friend promises to give us \$100 for the next three years starting next year, the present value of those cash flows assuming the term structure is as presented in table 6 (and our friend can be trusted to actually give us the money) is

$$\begin{aligned} Value_0 &= \frac{100}{(1+0.0040/2)^2)} + \frac{100}{(1+0.0078/2)^4)} + \frac{100}{(1+0.0104/2)^6)} \\ &= \frac{100}{(1+0.0040)} + \frac{100}{(1+0.007815)^2)} + \frac{100}{(1+0.010427)^3)} \\ &= \$295.35. \end{aligned}$$

A few comments about this calculation. First, we used different discount rates to discount the cash flows. This is what the term structure tells us - the appropriate discount rate for each cash flow. Second, because the interest rates in the term structure are quoted as APRs, we had to convert them to periodic rates *or* effective annual rates before discounting the cash flows. The equations above show both approaches. (Treasury securities, as we'll seen in chapter (7), have semi-annual compounding, which is why we divided by 2 above.)

The bottom line is that when the term structure is not flat, i.e., interest rates differ depending on the loan term, we either have to discount cash flows at different points in time by different discount rates, or we have to make a simplifying assumption to use the same discount rate to discount all cash flows. We'll see both approaches used throughout this book as they are in practice.

3.6.3 Different Interest Rates Do NOT Mean Better or Worse!

What the different interest rates implied by the term structure do *not* mean is that some loans or investments are necessarily "better" than others. Consider the mortgage rates from

Copyright © 2020 Michael R. Roberts

Table 4. It is tempting to think that because the 15-year mortgage has the lowest interest rate (2.4%) that it is the best mortgage from a financial perspective. It is not! Comparing the interest rate on a 15-year mortgage to a 30-year mortgage is like comparing the price of apples to the price of oranges. Any difference in prices between the two fruits tells us about differences in supply and demand across the two markets, but it doesn't tell us that one fruit is "better" than the other.

The lower interest rate on the 15-year mortgage just means that there is relatively less demand (more supply) for these mortgages than for the 30-year, for example. This is not surprising because 15-year mortgages come with significantly higher monthly payments, all else equal. Which financial product is "better" for us depends on several factors including our opportunity cost, our income, our opinion of where future interest rates may be headed, and how long we anticipate staying in the home.

The flip side of this story is that higher interest rates do not necessarily imply better investments. We discussed this in the first chapter when we noted that different returns were largely a consequence of different risks. Even if Treasury securities are guaranteed to make all of their payments, we'll see later that their value is still sensitive to changes in interest rates. In other words, should we want to sell a Treasury security before maturity, we are exposed to risk because the price of the security will depend on what happens to future interest rates - something we can't predict.

But, even putting risk aside, we can't look at the 2022 interest rates on Treasury securities in table 6, for example, and assume that the 20-year bond is the best Treasury investment because it offers the highest interest rate. If this were true, investors would start buying the 20-year Treasury bond, which would drive up the price. As we'll discuss in chapter 7, the increase in price leads to a decrease in the yield on the bond.

Another way to see the problem with this logic is to imagine buying the 20-year bond today. This guarantees you an average annual return of 2.05% per year for the next 20 years. But, what if interest rates jump up to 5% the day after you buy your 20-year Treasury. Now you're disappointed (to say the least) because you're stuck with a bond paying you 2.05% when new bonds will be paying 5%.

So, the punchline of this section is as follows. The term structure of interest rates tells us how the interest rates on loans differ according to their term. Because loans of different terms are different products - like apples and oranges - the different interest rates do *not* imply better or worse from a financial perspective, just different **equilibrium** outcomes in the supply and demand of those different products.

Copyright © 2020 Michael R. Roberts

3.6.4 The Fed and Monetary Policy

The **federal reserve system**, or **the Fed**, is the U.S. central bank.[12] This bank consists of the board of governors, the federal reserve banks, and the federal open market committee. These three entities are responsible for:

- Conducting monetary policy,
- Maintaining stability of the financial system,
- Supervising and regulating financial institutions, such as banks,
- 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 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. By spending less, the growth in the prices of goods and services (i.e., inflation) should slow down because of slowing demand. Hence, inflation is controlled by increasing interest rates.

However, raising interest rates increases r in our fundamental value relation (equation 1.2). Absent an offsetting increase in cash flows, values will fall as r increases; we're dividing by a bigger number. In other words, a consequence of the Fed raising interest rates - **monetary tightening** - is that asset values will fall. That's exactly what's happened since the Fed began its policy of raising interest rates. Between March and September of

[12] For more details, see the federal reserve website at: https://www.federalreserve.gov/default.htm.

Copyright © 2020 Michael R. Roberts

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 recession 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 valuation 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.

3.7 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.

- 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. They are linked by equations 3.1 and 3.3. The EAR and periodic interest rate are both discount rates that may be used to discount cash flows in our fundamental value relation. The APR is *not* a discount rate and 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.

Copyright © 2020 Michael R. Roberts

- Compound interest refers to interest earning interest and is a consequence of more frequent compounding (e.g., semi-annual, quarterly, monthly, daily).
- Mortgages, leases, and credit cards are all just loans with monthly interest accrual and monthly payments. The first two are amortizing loans because the payments are fixed over the maturity. The last is more flexible in terms of repayment size.
- The term structure of interest rates tells us the interest rates for loans of different lengths or terms. We'll see in future chapters that we can use these rates to discount cash flows arriving or departing at different points in time.
- The Fed's 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.

3.8 Technical Appendix

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

$$\lim_{k\to\infty}\left(1+\frac{APR}{k}\right)^{k} = e^{APR} \tag{3.6}$$

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

$$\lim_{k\to\infty}\left(1+\frac{APR}{k}\right)^{kT} = e^{APR\times T} \tag{3.7}$$

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 rarely continuously compounded in real life, assuming so makes some financial calculations easier. We won't use it much in this book, but it is important to be familiar with the concept.

Copyright © 2020 Michael R. Roberts

3.9 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:

a. annually,

b. quarterly,

c. monthly, and

d. 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:

a. What is the monthly periodic interest rate?

b. What is the effective annual interest rate (EAR)?

c. What is the monthly payment?

d. 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?

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

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 (*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:

Copyright © 2020 Michael R. Roberts

a. What is the principal amount of the loan?

b. What is the monthly mortgage payment?

c. How much money is owed on the loan after the 83rd payment?

d. What fraction of the 84th payment is interest? Principal reduction?

e. (*Challenging*) How long will it take to repay one half of the loan?

3.16 (*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?

3.17 (*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?

3.18 (*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?

Copyright © 2020 Michael R. Roberts

3.19 (*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?

3.20 (*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?

3.21 (*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.

Copyright © 2020 Michael R. Roberts

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.

3.22 (*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?

3.23 (*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.

Copyright © 2020 Michael R. Roberts

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?

3.24 (*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?

Copyright © 2020 Michael R. Roberts

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.)

3.25 (*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 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?

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Part II

Decisions Most People Make

Chapter 4

A Financial Perspective of Business

Fundamental value relation

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

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. That said, the lingo can be daunting and the settings can seem 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, and

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.

Copyright © 2020 Michael R. Roberts

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 (surface laptops and tablets, Xbox consoles), and cloud services (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.

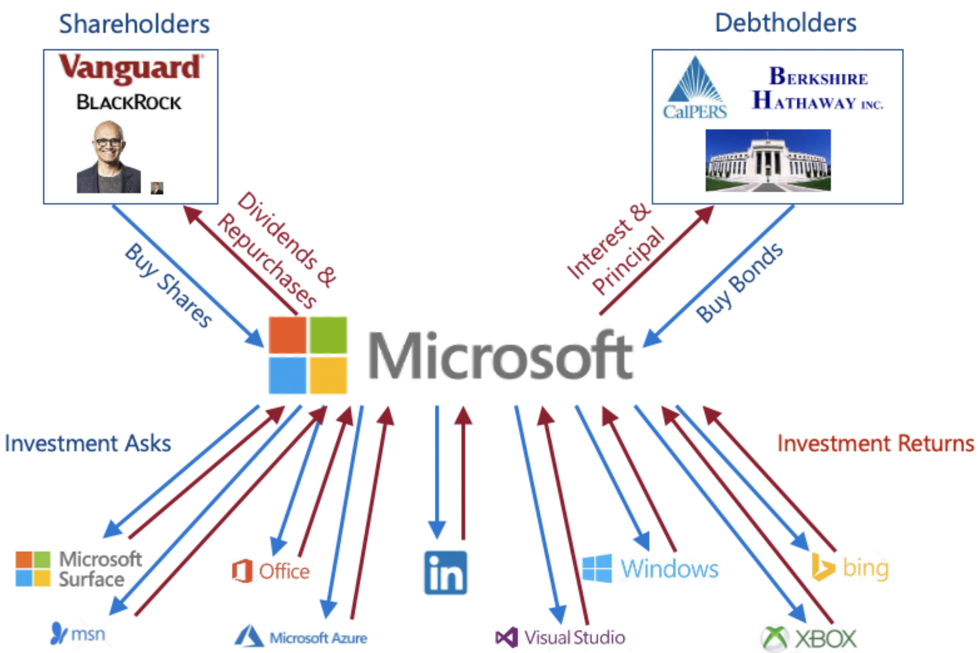

Figure 4.1: Microsoft from a Financial Perspective

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. We'll discuss later what it means to be a debtholder (chapter 7) or shareholder (chapter 8). For now, it suffices to understand that these investors are responsible for giving resources 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

Copyright © 2020 Michael R. Roberts

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 in investor accounts.

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 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 logos for CalPERS and Berkshire Hathaway. CalPERS, which stands for the California

[1] To 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.

Copyright © 2020 Michael R. Roberts

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 compute stock prices, thereby determining how much they should pay to purchase shares. 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 return on the stock, what investors expect to earn on each dollar they invest. 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 return on the loan.

Notice that the cash flows and expected returns to investors are costs to companies. What investors expect to earn is what companies are expected to pay. Hence, r in the fundamental value relation is referred to as a **cost of capital** for companies. The cost of capital measures how much companies have to pay their investors for each dollar they

[2] Much 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.

Copyright © 2020 Michael R. Roberts

receive. So, expected returns to investors and costs of capital to companies are two sides of the same coin. Investors earn returns that are paid by companies.

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.

These investment returns are indicated by the red arrows pointing from the different projects back to Microsoft. It’s at this point, the business has another decision: What to do with the investment returns. As discussed above, 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.

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.

Copyright © 2020 Michael R. Roberts

Critically, we'll see that companies, just like us, should use the NPV criterion to make their decisions.

4.3 Corporate Social Responsibility

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) initiatives as evidence that companies are concerned about more than just profits. CSR refers to businesses making societal concerns an important consideration in their decision-making. 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 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.

Copyright © 2020 Michael R. Roberts

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.

- 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?
- How do companies determine where to invest their money - shareholders or debtholders - and does it matter?

The remainder of the book is aimed largely 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 while in retirement? That return comes from the companies in which we invest. We give them money. This is our investment. Companies invest our money. If those investments pay off, the companies return to us even more money than what we gave them. So, the return we as savers earn on our investments in companies is the cost of raising money faced by companies. These are two sides of the same coin and therefore are referring to the same r in our fundamental valuation relation.

[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.

Copyright © 2020 Michael R. Roberts

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 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.

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).

Copyright © 2020 Michael R. Roberts

Chapter 5

Project Viability

Fundamental value relation

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

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 (we'll recognize the most important one),
- show how the cash flows from our fundamental value relation are computed in a specific way that is applicable to all business decisions,
- explicitly link common key performance indicators (KPIs) on which managers focus (revenue growth, margins, earnings, days receivable/payable, inventory turnover, etc.) to cash flow and value,
- apply our fundamental value relation to investigate whether Dell Inc. should produce and sell a tablet, and
- highlight the advantages and limitations of business decision making by performing sensitivity analysis on our tablet project.

Copyright © 2020 Michael R. Roberts

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. This chapter assumes you are comfortable with basic accounting. If not, read or refer back to appendix A.

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 multiple examples beginning with a very simple, stylized example in which we have a sales initiative costing \$100 million today and generating \$50 million 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.

Years	0	1	2	3
Cash flows (\$mil)	-100	50	50	50

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, in this example that's the \$100 million up-front cost and the \$50 million each year thereafter. For example, if the company's existing sales are projected at \$2 billion per year over the next three years regardless of whether the sales initiative is undertaken, then these existing sales are irrelevant for the evaluation of the initiative. If instead, the sales initiative will also boost existing sales by 1% for the next three years, in addition to generating \$50 million per year, we need to consider that boost (\$20 million) in the evaluation of the sales initiative.

To determine what's incremental, we need to identify what *changes* as a result of the project. 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. We'll see further examples below.

[1] Graham and Harvey (2001) show that these three criteria are the most popular decision criteria used by Fortune 500 CFOs.

Copyright © 2020 Michael R. Roberts

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.

$$NPV_0 = CashFlow_0 + \frac{CashFlow_1}{(1+r)} + \frac{CashFlow_2}{(1+r)^2} + \frac{CashFlow_3}{(1+r)^3} + ... \tag{5.1}$$

Often the initial cash flow, $CashFlow_0$, is negative and corresponds to the cost or price of the investment. However, there is no limitation on the signs of the cash flows. An investment may have many outflows that can occur at any points in time.

If a manager announced to the stock market the NPV of a project, and the market believed them, the value of the company's equity, its **market capitalization**, would change by the NPV of the project.[2] As such, the **NPV decision criteria** says accept a project when the NPV is positive, reject when negative.

The NPV ($mil) of our sales initiative is

$$NPV_0 = -100 + \frac{50}{(1+0.10)} + \frac{50}{(1+0.10)^2} + \frac{50}{(1+0.10)^3} = \$24.34 \text{ million.}$$

Thus, we would accept this project because it creates over $24 million of value for the owners of the company.

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

$$Value_0 = \frac{50}{(1+0.10)} + \frac{50}{(1+0.10)^2} + \frac{50}{(1+0.10)^3} = \$124.34 \text{ million.}$$

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 that creates $24.34 million in today's dollars.

From where does this value come? Competitive and comparative advantage, technological innovation, barriers to entry, 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 creation when positive, value destruction when negative.

Before turning to our other business decision criteria, it's important to emphasize the similarity between the NPV criterion used in business and that used in personal finance. Simply put, they're the same! In our personal lives, we strive to make positive NPV decisions. Businesses do the same.

[2] Market capitalization is the number of shares outstanding times the price per share.

Copyright © 2020 Michael R. Roberts

5.1.2 Internal Rate of Return (IRR)

The **internal rate of return** (**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.

$$NPV_0 = 0 = CashFlow_0 + \frac{CashFlow_1}{(1+IRR)} + \frac{CashFlow_2}{(1+IRR)^2} + \frac{CashFlow_3}{(1+IRR)^3} + ... \tag{5.2}$$

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.[3]

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).

$$Price_0 = \frac{CashFlow_1}{(1+IRR)} + \frac{CashFlow_2}{(1+IRR)^2} + \frac{CashFlow_3}{(1+IRR)^3} + ...$$

Viewed this way, the IRR is the periodic return investors earn from the project.

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, the cost of capital, and invest it at a higher rate IRR, the internal rate of return, thereby creating value.

Consider our sales initiative. Using a computer, we can solve the following equation to get the IRR.[4]

$$NPV_0 = 0 = -100 + \frac{50}{(1+IRR)} + \frac{50}{(1+IRR)^2} + \frac{50}{(1+IRR)^3} \implies r = 0.2338$$

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.

[3]ROI 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.

[4]In Excel, the function *IRR* will return the IRR for a set of cash flows.

Copyright © 2020 Michael R. Roberts

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. That's a good thing. So, we would accept this project on the basis of the IRR criterion. Though, as we'll see in chapter 6, the IRR suffers from several shortcomings that limit its use.

The threshold return that projects' IRRs 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 setting an artificially high hurdle rate for projects to clear is an increase in the risk of the business.

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 1.

	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 1: 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. In period 1, cumulative cash flow is the previous period's cumulative

Copyright © 2020 Michael R. Roberts

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.

A slight variant on the payback period is the **discounted payback period**, which is computed in the same manner only using discounted cash flows. The bottom two rows of the table 1 present the discounted and cumulative discounted cash flows. The discounted payback period is a little over two years.

The payback period is an appealing decision criterion 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 biggest is that the payback period has nothing to say about the value of a project. Another shortcoming is that the payback period criterion can lead to myopic decision making. Decision makers will select projects that payoff more quickly, as opposed to projects that are more valuable. Finally, the threshold for what distinguishes acceptable from unacceptable projects is largely arbitrary and removed from value considerations. Most implementations of the payback period rely on historical performance of similar projects for guidance on what is and what is not an acceptable payback period. Thus, while informative, the payback period has some serious drawbacks.

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, rather than getting lost in semantics, let's understand what they are measuring and their shortcomings.

Cash-on-cash return or **cash yield** is computed as the ratio of the periodic cash received to the total cash invested. Some use a pre-tax measure of cash flow, others use a post-tax measure. Ultimately, the cash-on-cash return tells us the amount of money generated per

Copyright © 2020 Michael R. Roberts

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 $(50 + 50 + 50)/100 = 1.5$x. 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, hence the time value of money. Table 2 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 and risk profiles, and even then is potentially problematic.

	Cash Flows ($mil)					
Projects	0	1	2	3	NPV (r=10%)	Cash Multiple
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 2: Cash Multiple vs. NPV

It's natural to ask why practitioners use some of these metrics despite known shortcomings? Some practitioners 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. While none of these reasons are a justification for their use, they do help rationalize it.

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.

Copyright © 2020 Michael R. Roberts

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 called **Discounted Cash Flow Analysis**, or **DCF** for short.

Many academics will say: "Use the NPV criterion and forget the rest." Strictly speaking, they're correct. Practically speaking, all three 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. Finally, 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 by itself. We'll see examples 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 everyone, regardless of your particular function within the company. If you're going to justify spending money on anything, you must present the financial costs and benefits, i.e., the cash flows.

The cost of capital is the purview of finance specialists, such as the CFO. Financial executives should tell you what the cost of capital or hurdle rate is *before* you present your business case. In other words, you 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 purchases, cash flows are referred to as **free cash flows**. 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.

Copyright © 2020 Michael R. Roberts

Free cash flow is the money generated or needed by a project after all revenues, expenses, investments, and taxes have been considered. When positive, companies have extra money that they can do with as they please, such as

- distribute to creditors by paying interest or reducing 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 that company will be, all else equal, according to our fundamental value equation. (The numerator of each term, $CashFlow$, is bigger.)

In order to generate positive free cash flows, companies often need to generate negative free cash flows. Put differently, to make money, you have to spend money. When free cash flow is negative, companies must find money to fund the project. 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

$$\begin{aligned}\text{Unlevered free cash flow} &= (Sales - Expenses - D\&A) \times (1 - \tau) \\ &\quad + D\&A - NLTI - NWCI. \end{aligned} \tag{5.3}$$

That equation looks daunting, but is in fact a representation of how businesses function. Let's define the acronyms. $D\&A$ stands for depreciation and amortization, τ is the marginal tax rate, $NLTI$ stands for net long-term investment, and $NWCI$ stands for net working capital investment. We discuss each of these components and how they affect cash flow below.

Copyright © 2020 Michael R. Roberts

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**). It seems every financial analyst came up with their own acronym.

NOPAT starts with *Sales*. While many projects are aimed at generating sales, 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 value in these examples.

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, 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 ($D\&A$).[5] *Sales* minus *Expenses* and $D\&A$ is **EBIT** - Earnings Before Interest and Taxes.

Finally, $EBIT$ times $(1 - \tau)$ give us what we started with, NOPAT. Let's annotate our free cash flow expression (equation 5.3) with the different earnings measures.

$$\begin{aligned}\text{Unlevered free cash flow} &= \underbrace{\overbrace{\left(\overbrace{Sales - Expenses}^{\text{EBITDA}} - D\&A\right)}^{\text{EBIT}} \times (1 - \tau)}_{\text{NOPAT}} \\ &\quad + D\&A - NLTI - NWCI.\end{aligned}$$

Free cash flow starts with after-tax operating profits, which is built up from measures of operating earnings (EBITDA and EBIT).

[5]Strictly 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. For most practical capital budgeting applications, this level of detail is unnecessary.

Copyright © 2020 Michael R. Roberts

Notice we did not subtract interest expense in our NOPAT calculation. In other words, NOPAT is similar to net income from the income statement except NOPAT excludes interest expense (and income) and taxes are computed off of EBIT as opposed to pre-tax income. How much interest the company pays on what it borrows does *not* affect unlevered free cash flow. Thus, the NOPAT and net income of any project will differ if the company borrows money to pay for the project.

By not subtracting interest expense, we ensure free cash flow is unaffected by how the project is funded, with debt or equity. Where our CFO gets money to pay for our project - loans, cash savings, stock issuances - is irrelevant for our calculation of unlevered free cash flow. This is important because it leaves decision makers free 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, that is raise money with the lowest cost of capital, r. Everyone else's job is to find projects with really big free cash flows.

To see why, recall the fundamental valuation relation (equation 1.2). Value is driven by (i) cash flows in the numerators and (ii) discount rates in the denominators. Everybody in the company seeks to find projects with big (free) cash flows. The CFO seeks to find money with a really low cost of capital. The result is a nice separation of tasks and a project with a lot of value.

5.2.2 Investments

In addition to ignoring interest expense, free cash flow differs from profits 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. No money leaves the company when an asset is depreciated or amortized.

You might wonder: 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

Copyright © 2020 Michael R. Roberts

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 tax rate on the project's profits. 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.

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.

The final term is **net working capital investment** ($NWCI$), also known as the **change in net working capital** or sometimes just the **change in working capital**. Recall that net working capital is the difference between current assets (e.g., inventory, accounts receivable) and current liabilities (e.g., accounts payable). 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. That is, the company spent \$200 on its net working capital perhaps building up inventory, extending credit to its customers, or paying its suppliers. 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 afloat on a day-to-day basis.

Two subtle points regarding net working capital investment. First, we exclude any financing. This means ignoring short-term debt and long-term debt due within a year, both of which are current liabilities. Remember, unlevered 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.

Copyright © 2020 Michael R. Roberts

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 affected.

5.3 Application: Choosing a Stand Alone Project

The best way to understand capital budgeting is to do it. We'll start by considering a stand alone project, that is, whether to invest in a project independent of other projects. The following example is hypothetical.

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

Copyright © 2020 Michael R. Roberts

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 improvements in the product.
    - 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 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 construct the production facility, which has 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 liquidation value to move any future production offshore.
4. $NWCI$, Net working capital investment

Copyright © 2020 Michael R. Roberts

- 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.

Let's assume that Dell's marginal tax rate is 21%, and the cost of capital for this 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 top row identifies the period (year) of the project's life. As always, period 0 is today, January of 2012.

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)

Copyright © 2020 Michael R. Roberts

5.3.2 Incremental Earnings

Figure 5.3 presents the **pro forma**, or forecasted, P&L (a.k.a., profit and loss statement) for the project. This statement shows incremental sales, expenses, taxes, and earnings for the project.

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 5.3: Dell Tablet - P&L

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.

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 sales. 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 differently, **capital budgeting exercises must capture *all* of the incremental effects associated with the project - direct and indirect**.

In our tablet project, the reduction of laptop sales only happens because we are introducing the tablet. Therefore, it must be accounted for in our DCF analysis. In contrast,

Copyright © 2020 Michael R. Roberts

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 there is no change and 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. 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. R&D expenses are given.

Depreciation and amortization are a consequence of the capital expenditures on the production facility. The $400 million facility 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. This tax rate is the **marginal tax rate** of the company. In other words, it is the tax rate applied to each additional dollar of income. This is different from the **effective tax rate** which is a blend of the different tax rates applied to the company's earnings and is more relevant for valuing the entire company as opposed to an individual project.

Notice taxes in the first two years are positive. 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 occured *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

Copyright © 2020 Michael R. Roberts

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

Had Dell not been profitable or the tablet project been the only project at Dell, there would be no taxable income to shield. In this case, Dell could use **operating loss carryforwards** to shield *future* profits from taxes. This analysis is displayed in figure 5.4.

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

In period 0, the $50 million operating loss 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 wiped out by the $187.9 million dollars of operating income in period 2, but our taxable income is significantly reduced. The impact on the P&L can be seen in figure 5.5.

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, whose calculation was detailed above.

Dell spends $400 million in capital expenditures today to develop and construct the facility. Three years later Dell will receive a $200 million cash inflow from the sale of the

Copyright © 2020 Michael R. Roberts

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

facility - its liquidation value. However, this $200 million ignores taxes on any **capital gains** (i.e., profits from the sale of an asset), which are computed as

$$\text{Tax on capital gains} = \underbrace{(\text{Sale Price} - \text{Book value})}_{\text{Capital gains(losses)}} \times \text{Tax rate}$$

The book value of the asset is the purchase price or original cost ($400 million) less the accumulated depreciation. Because the asset will have been in use for three years at the time of sale, the book value is

$$\text{Book value} = \text{Purchase price} - \text{Accumulated depreciation} = 400 - (3 \times 39) = \$283 \text{ million}$$

The book value is tracked annually in figure 5.6 as the net PP&E, where net refers to net of accumulated depreciation.

If we assume the tax rate on capital gains is the same as the corporate income tax rate (21%) then

$$\text{Tax on Capital Gains} = (200 - 283) \times 0.21 = -\$17.43 \text{ million.}$$

The negative taxes mean Dell is anticipating a tax shield from the loss it will incur on the sale of its production facility.

Copyright © 2020 Michael R. Roberts

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

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.

Accounts receivable are computed as

$$\text{Accounts receivable} = \underbrace{\frac{\text{Days receivable}}{360}}_{\text{\% Sales Outstanding}} \times \text{Sales}. \tag{5.5}$$

There is a nice intuition to this calculation. Dell takes 14 days to collect on sales to its customers. If sales are distributed uniformly throughout the year, then at any point in time $14/360 = 0.0389$, or 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 is often used to ensure equal length years and 90-day quarters, but it doesn't really matter which you use because the difference is small.)

Copyright © 2020 Michael R. Roberts

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

For example, after one year, Dell will have made $75 million in net sales. At the end of that year, 3.9% of this $75 million in sales, $2.92 million, is still uncollected and show up as accounts receivable on their balance sheet. In the second year of the project, 3.9% of the $1.09 billion in sales, $42.4 million, is still outstanding at the end of the year.

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 mean Dell stars year one with 0.2 million units in their inventory, costing them $3.75 million. As time progresses, Dell sell's 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

Copyright © 2020 Michael R. Roberts

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 (e.g., $49.3-3.8 = 45.6$, $197.3-49.3 = 148$). Remember, balance sheet accounts measure the *stock* of value at a point in time. We want the *flow* of money over each period, which requires computing the change in these accounts.

For example, the $2.9 million increase in receivables from period 0 to 1 is money Dell is not collecting and therefore a cash outflow. Likewise, the buildup of $3.8 million of inventory in period 0 corresponds to a cash outflow. In contrast, the increase in accounts payables corresponds to money Dell is not paying and therefore cash inflows. Yes, Dell owes the money but only later. Therefore, it can invest the money today.

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 cahs 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. Strictly speaking, Dell will collect the $168.6 million from their customers during the first two weeks of the fourth year. Likewise, they'll pay their suppliers

Copyright © 2020 Michael R. Roberts

$89.5 million during the first 45 days of the fourth year. Rather than create another column in our spreadsheet, we've just assumed that these collections and payments all occur at the end of the third 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 unit so close (3 versus 3.2), it's a small enough sin to ignore.

5.3.5 Free Cash Flows and Valuation

Figure 5.8 puts everything together to compute the project's free cash flows.

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

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 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 assumed 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

Copyright © 2020 Michael R. Roberts

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. The point of 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.

Finally, 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* (**KPI**s) 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

Copyright © 2020 Michael R. Roberts

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 the organization 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 tells 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 above*? Table 3 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

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 3: KPI Response to Lower Market Sales (\$mil)

Copyright © 2020 Michael R. Roberts

million; IRR falls from 39.2% to 16.49%. The payback period gets pushed out but only modestly - a little over two months. Year one earnings losses jump in from $29.2 to $36.0 million. 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

		Market Share						
		10%	15%	20%	25%	30%	35%	40%
Unit Price ($)	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

corresponds to a different initial unit price point, 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 corner of the matrix, at the highest NPV, 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.

Copyright © 2020 Michael R. Roberts

5.4.2 Cost Assumptions

- *What if inflation leads to faster unit cost increases of 7% per annum, as opposed to 3%?* Again, the qualitative implications of an increase in input costs are obvious;

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

  value will decline all else equal. By how much is less clear. The table shows that project value is only modestly affected, as is liquidity. 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 - fixed costs-to-variable costs - for the project decrease because variable costs increase relative to overhead. That is, the operating risk of the project decreases.

- *There is concern that R&D 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 is largely unaffected, though first

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

  year earnings 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 value most sensitive, i.e., the key value drivers?

Copyright © 2020 Michael R. Roberts

Table 4 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.

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 4: Dell Table: Key Value Drivers

$$\frac{\text{Change in NPV}}{\text{Change in unit price}} \times \frac{\text{Unit price}}{\text{NPV}} = \frac{430.2 - 417.5}{404 - 400} \times \frac{400}{417.5} = 3.06$$

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

Copyright © 2020 Michael R. Roberts

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 reality, this may not be possible. For example, changing the sales 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 "large" 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 5 presents the baseline and break-even values for each forecast driver. The break-

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 5: Dell Table: Baseline and Break-Even Values

even values must be calculated numerically (I used GoalSeek in Excel). For example, if unit costs increase to $131.60, the NPV of the project will be zero, anything higher, the NPV is negative, and the project is 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 con-

Copyright © 2020 Michael R. Roberts

trol 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 actually pretty funny, so days payable is obviously not a risk factor. 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 6 presents the baseline and target values for a subset of forecast drivers. At the target

	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 6: Dell Tablet: Targeting the IRR

value for each forecast driver, the IRR of the project equals 45%. (I used GoalSeek in Excel to find these values.) So, 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.

- *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

Copyright © 2020 Michael R. Roberts

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 is $0.20 \times \$75.8 = \$15.16$ million.

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 name changes, the mechanics and intuition are identical to personal financial decisions.

- 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.

Copyright © 2020 Michael R. Roberts

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.

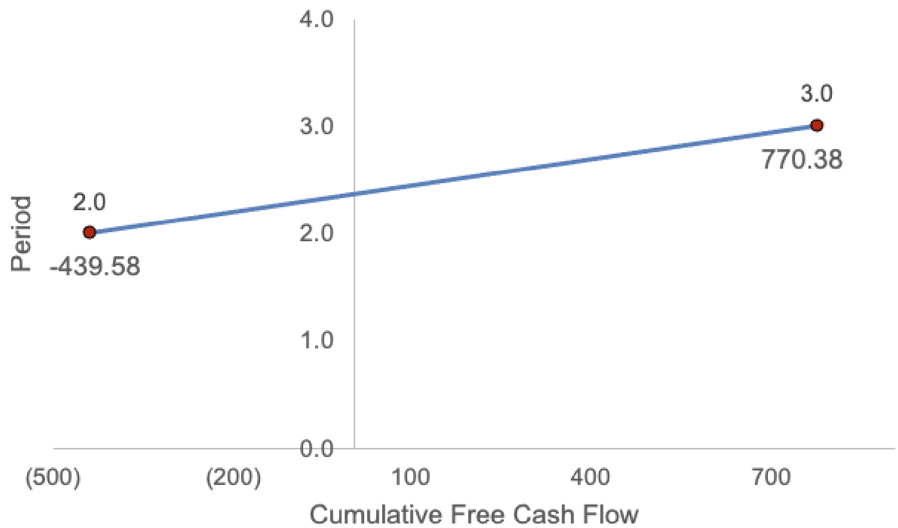

Figure 5.11: Linear Interpolation of Cumulative Free Cash Flows from Dell Tablet Project

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

$$\text{Period} = 0.0008264736024 \times \text{Cumulative Free Cash Flow} + 2.3633012662.$$

Copyright © 2020 Michael R. Roberts

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.

$$y = y_0 + (x - x_0)\frac{y_1 - y_0}{x_1 - x_0}$$

Copyright © 2020 Michael R. Roberts

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.

Free Cash Flow

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 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

Copyright © 2020 Michael R. Roberts

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 forecasted 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.

Copyright © 2020 Michael R. Roberts

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 analsis 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.

Copyright © 2020 Michael R. Roberts

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 forecasted 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.

Copyright © 2020 Michael R. Roberts

	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

(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 depreciation before the asset is sold or retired.) The different depreciation schedules are illustated in Figure 5.12.

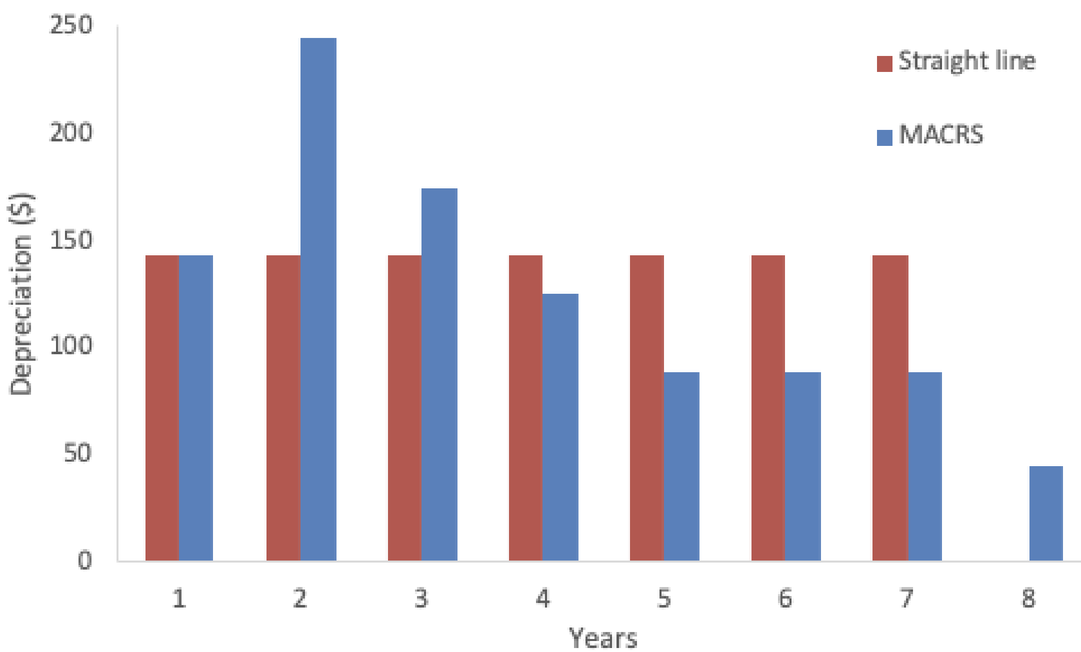

Figure 5.12: Straight Line versus MACRS Depreciation

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?

Stand Alone Investments

5.12 (*DCF, decision criteria*) For a project with the following cash flows, calculate the

- Net present value,

Copyright © 2020 Michael R. Roberts

- Internal rate of return,
- Payback period,
- Discounted payback period,
- Profitability index, 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.)

Copyright © 2020 Michael R. Roberts

Using this information, answer the following questions.

a. What is the pro forma P&L statement for the lemonade business? Specifically, what are the revenues, expenses, and (before- and after-tax) profits?

b. What are the free cash flows for the business?

c. 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.

a. What are the NPV, IRR, and Payback period for this investment opportunity?

b. Based on your answers to the previous question, would you invest in this opportunity? Explain your answer.

c. What would your brother have to pay you in four years for you to break even?

d. 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.

a. How much revenue will Jerry generate each year?

Copyright © 2020 Michael R. Roberts

b. What are Jerry's before- and after-tax earnings each year?

c. What is the value of Jerry's newsletter in today's dollars?

d. (*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 forecasted 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.

a. What are the NPV and IRR of Hooker Chemical's new plant?

b. Based on your answers to the previous question, should Hooker invest in the new plant? Explain your answer.

c. What is the break-even annual clean-up fee?

d. 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*:

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

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 Phillie's 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.

Copyright © 2020 Michael R. Roberts

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.

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

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 $15,000,000 \times (1 + 0.04) = \$16,500,000$. 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

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

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 salvage value 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, and their cost of capital is 15%.

Using this information, answer the following questions.

Copyright © 2020 Michael R. Roberts

- 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?

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Chapter 6

Project Selection

Fundamental value relation

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

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.
- examine the implications of comparing projects with different lifetimes,
- demystify 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 large majority of corporate financial decisions. Two other corporate decisions - financial policy and acquisitions - are more specialized and presented in the third part of the book.

Copyright © 2020 Michael R. Roberts

6.1 Application: Choosing Among Mutually Exclusive Projects

6.1.1 The IRR Rule and Cash Flow Signs

While frequently used for decision making, IRR has shortcomings. For assessing stand alone projects, the IRR criterion leads to the same decision as the NPV criterion only when all cash outflows precede all cash inflows. In other words, all negative cash flows must precede all positive cash flows. We'll call this the "sign rule" because the signs of the cash flows must abide by the rule that all negatives precede all positives. 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).

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

Our Dell tablet example adheres to the sign rule and, as such, the IRR and NPV criteria are in agreement on accepting the project. However, when choosing among projects, relying on IRR to identify the most valuable project can lead you astray even if all of the projects adhere to the sign rule. The hypothetical case study examined in this section illustrates these, and other pitfalls of the IRR criterion.

6.1.2 Vanguard's IT Overhaul

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 bidder provided two separate bids, all of which are detailed in table 2. Our job is to assess these bids assuming Vanguard's cost of capital is 12% and decide which one is best.

Copyright © 2020 Michael R. Roberts

Bidder	**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 2: Vanguard IT Bids

6.1.3 Using the IRR to Compare Projects

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 (\$000)	−100	60	60	60

Figure 6.1: Cisco 1 Bid Timeline

The NPV and IRR for Cisco 1 are

$$NPV_0 = -100 + \frac{60}{(1+0.12)} + \frac{60}{(1+0.12)^2} + \frac{60}{(1+0.12)^3} = \$44.11 \text{ million, and}$$

$$0 = -100 + \frac{60}{(1+IRR)} + \frac{60}{(1+IRR)^2} + \frac{60}{(1+IRR)^3} \implies IRR = 36.31\%.$$

NPV is positive. IRR is greater than the cost of capital. Thus, both decision criteria imply Cisco 1 is a viable bid. Because the cash flows adhere to the sign rule - all cash outflows before cash inflows - the NPV and IRR rules necessarily agree.

Copyright © 2020 Michael R. Roberts

Now consider the Cisco 2 bid whose timeline is in figure 6.2

Years	0	1	2	3
Savings ($000)		60	60	60
Costs ($000)	−20	−35	−35	−35
Cash flows ($000)	−20	25	25	25

Figure 6.2: Cisco 2 Bid Timeline

The NPV and IRR for Cisco 2 are

$$NPV_0 = -20 + \frac{25}{(1+0.12)} + \frac{25}{(1+0.12)^2} + \frac{25}{(1+0.12)^3} = \$40.05 \text{ million, and}$$

$$0 = -20 + \frac{25}{(1+IRR)} + \frac{25}{(1+IRR)^2} + \frac{25}{(1+IRR)^3} \implies IRR = 111.85\%.$$

Similar to Cisco 1, the NPV is positive and the IRR is greater than the cost of capital. Therefore, Cisco 2 is also a viable bid by both NPV and IRR decision criteria.

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, ultimately 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 should be clearly preferable because we get more money, $1 versus $10,000.

A closer look at Cisco 2 reveals another perspective on why the NPV of Cisco 1 is larger. Instead of paying $100 million today as in Cisco 1, Vanguard only pays $20 million today, and then $35 million over the next three years to receive the same cost savings. But, this is just a loan in which Cisco lends Vanguard $80 million today in return for the three $35 million repayments. The interest rate, r, on this loan is

$$80 = \frac{35}{(1+r)} + \frac{35}{(1+r)^2} + \frac{35}{(1+r)^3} \implies r = 14.93\%.$$

Copyright © 2020 Michael R. Roberts

Note, r=14.93% is just the internal rate of return, which I 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. 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.

Some 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%. If some are 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, you may be right but you're changing the rules of the game. You're 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.

6.1.4 Further Shortcomings of the IRR Rule

The cash flows, NPVs, and IRRs for all of the bids are presented in Table 3. Let's focus on understanding the remaining four bids. Juniper 1 has the same sized investment as Cisco 1, $100 million, and a higher IRR, but lower NPV. This comparison illustrates how differences in the timing of cash flows can create a wedge between the NPV and IRR criteria when comparing projects. Front loading the cash flows ramps up the return but at the expense of value in this comparison.

Copyright © 2020 Michael R. Roberts

	Period					
Bids	0	1	2	3	NPV	IRR
Cisco 1	(100.0)	60.0	60.0	60.0	44.11	36.31%
Cisco 2	(20.0)	25.0	25.0	25.0	40.05	111.85%
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	NaN
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: Vanguard IT Bid Results

Juniper 2 is interesting for several reasons. The cash flows do not adhere to the sign rule; an inflow precedes an outflow. There are also two sign changes in this cash flow stream, positive to negative going from period 0 to 1 and negative to positive going from period 1 to 2. Accordingly, there are actually two different IRRs. In fact, there are as many IRRs as there are sign changes in the cash flow streams. The problem this creates is that the IRRs are all different, but all equally possible. In other words, we can't simply select one IRR as being "better" than the others.

The case of Juniper 2 is particularly pernicious because the two IRRs are imaginary and not in the Harry Potter sense. They are complex numbers that are a function of the square root of negative one. If you don't understand this, that's fine because neither will anyone else (other than a mathematician). More practically, you do not want to justify a project by noting that it's return on investment is imaginary. Because Excel doesn't know what do in this situation, it spits out an error code (#NUM!). Yet, Juniper 2 is the most valuable bid according to the NPV criterion.

Huawei 1 is just another version of Cisco 2 - a small investment because of an embedded loan. Comparing the cash flows to Cisco 1, we see that the loan in Huawei 1 is for \$80 million with repayments of \$40 million each year thereafter. The implied interest rate is

$$80 = \frac{40}{(1+r)} + \frac{40}{(1+r)^2} + \frac{40}{(1+r)^3} \implies r = 23.38\%,$$

almost double Vanguard's cost of capital fo 12%.

Huawei 2 is more interesting. It has the second largest NPV (\$45.53 million) and a *negative* IRR. The reason for this disconnect is that all cash inflows precede all cash outflows. This is a loan. We get money in exchange for paying it back later. It may be a slightly weird loan in that we get money for several years before having to pay it all back in one lump

Copyright © 2020 Michael R. Roberts

sum, but it's still a loan. So, Huawei 2's cash flows are as if Huawei were lending Vanguard $50 million a year for three years starting today, and asking them to pay $125 million three years from today. The only thing better than a zero interest rate loan, is a negative interest rate loan in which the lender is paying us to take their money. This is what Huawei 2 is effectively doing, hence the positive NPV.

If you're wondering why Huawei would do this, don't forget that these cash flows reflect the cost and benefits experienced by Vanguard. So, while Huawei may only receive the $125 million three years from now, their costs to provide products and services are likely well below the benefits experienced by Vanguard (i.e., $50 million per year). So, this deal is likely positive NPV for Huawei, otherwise they would not make this bid.

The bottom line is that Juniper 2 offers the best bid from a value standpoint because it offers the highest NPV, assuming the bid has no affect on Vanguard's cost of capital or other investment opportunities.

6.1.5 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 do not make much sense or in conflict with NPV.

6.2 Application: Choosing Among Projects when Facing Constraints

Cypress Technologies is a hypothetical company that creates digital devices, such as thermostats, pressure gauges, and monitoring sensors. The Chief Marketing Officer (CMO) is tasked with providing marketing support for 6 different products, whose costs and benefits are detailed in Table 4.

The initial cost corresponds to the upfront cost of providing support to each product. The NPV is the estimated net present value of marketing support for the product. The challenge

Copyright © 2020 Michael R. Roberts

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 4: Cypress Technologies Product Lineup ($000s)

the CMO faces is that her budget is $600,000, less than the total cost of supporting all six products. Thus, the question is: Which products should marketing support to generate the most value? Equivalently, what is the best allocation of her limited financial resources?

6.2.1 Profitability Index

One approach to answering this question relies on the **profitability index** defined as

$$\text{Profitability Index} = \frac{\text{Net Present Value}}{\text{Amount of Resource Consumed}}. \tag{6.1}$$

The profitability index takes the value of a project as measured by NPV, and scales it by the amount of the constraining resource used in the project. In other words, the profitability index measures the value per unit of resource.

In our Cypress Technologies example, the constraining resource is money; the CMO doesn't have enough money to fund all of her projects. Table 5 presents the profitability index for each project. We have also sorted the projects by their profitability indices from highest to lowest, and computed the cumulative cost of the projects moving down the table.

The most valuable project is PressureView with an NPV of $1.365 million. However, the most valuable project per dollar spent is PressureView+, with a profitability index of 5.0. For each dollar spent on PressureView+, Cypress generates $5 of value in today's dollars. The CMO can now allocate her limited budget to projects offering the biggest bang for the buck until she exhausts all of her money. Table 5 shows that she should undertake all projects above the dashed line (PressureView, PressureView+, Thermo Chem, and Thermo). These projects offer the most value for her money, and completely exhaust her budget of $600,000. If she wanted to undertake the SensorView and Thermo Rx projects, she needs to get more money.

Copyright © 2020 Michael R. Roberts

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 5: Cypress Technologies Product Profitability Indices ($000s)

What happens if the constraining resource isn't money? Imagine that our CMO has a large enough budget to undertake all six projects, but she only has 20 direct reports that can work on these projects? Table 6 provides the relevant information for this situation.

Product	Headcount	NPV	Profitability Index	Cumulative 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 6: Cypress Technologies Product Profitability Indices ($000s)

The headcount column shows how many marketers 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 person, $123,600. The cumulative headcount column provides a running total of how many marketers are used by the projects.

The decision making process here is the same as before. We've ranked projects by their profitability indices in Table 6 and selected the highest ranking projects until we exhaust our resource. In this case, the CMO can undertake the SensorView, Thermo, and PressureView projects with her current headcount. If she wants to undertake any other projects, she's going to need to hire more people.

A nice feature of the profitability index in this case is that it tells us how much we can pay these new hires and ensure we are still making money. More precisely, the profitability index in this case tells us the most we can spend on an additional employee on an after-tax

Copyright © 2020 Michael R. Roberts

present value basis. For example, if we can hire two more marketers whose average after-tax cost over the course of the PressureView+ project has a present value less than $97,000 per marketer, then we can hire them, undertake the PressureView+ project, and still create value for the company.

The **profitability index decision rule** is as follows.

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.

Notice that the CMO will have $20 - 17 = 3$ extra employees that would seem to be doing nothing. There are a couple of options here. One is to fire them, though that can be very costly. Two is to allocate them to the accepted projects, though that assumes doing so does not destroy value (i.e., reduce NPV). It also begs the question of why they weren't allocated to those projects in the first place if they could create additional value for the projects.

This extra headcount can also create a problem for our decision rule. Imagine there exists a seventh project called ExoScale that requires 3 marketers and has an NPV of $40,000. The profitability index of this project is $40/3 = 13.3$, the lowest of all the projects listed in Table 6. However, rather than having the three extra marketers sit around doing nothing, we might allocate them to ExoScale. But, doing so violates our decision rule because there are other projects with higher profitability indices. This example highlights that the profitability index decision rule requires us to completely exhaust the resource. Otherwise, we may want to select projects with lower profitability indices to deploy some unutilized resources.

What if our CMO faced multiple constraints, for example, both headcount and budget? Now the problem becomes trickier and requires more advanced mathematical techniques, such as integer and linear programming. There is software that makes solving this problem relatively simple, but beyond the scope of this book and, frankly, most practitioners. One practical, if imperfect, solution to this problem is to apply the profitability index decision rule using the most binding or important constraint.

Copyright © 2020 Michael R. Roberts

6.3 Application: 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.

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 7 details the cash flows for each machine over one lifetime.

Year	Machine A Cash Flows ($)	Machine B Cash Flows ($)
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 7: 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.

$$\text{Machine A NPV} = -800,000 + \frac{500,000}{0.10}\left(1-(1+0.10)^{-3}\right) + \frac{250,000}{(1+0.10)^4} = \$614,179.36$$

$$\text{Machine B NPV} = -1,000,000 + \frac{600,000}{0.10}\left(1-(1+0.10)^{-3}\right) = 492,111.19$$

Comparing these NPVs suggests that machine A is superior; it has a higher NPV. Unfortunately, this comparison is invalid, because the two machines have different lives - four versus

Copyright © 2020 Michael R. Roberts

three years - *and* they will need to be replaced by another machine of the same type. To compare the value of these two machines 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 the machine over a single period. In our Roarke example, cash flows come and go each year so one period equals one year. 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 2.6.

$$\begin{aligned}\text{Machine A annuity equivalent cash flow} &= \frac{614,179.36 \times 0.10}{(1-(1+0.10)^{-4})} = \$193,755.66 \\ \text{Machine B annuity equivalent cash flow} &= \frac{492,111.19 \times 0.10}{(1-(1+0.10)^{-3})} = \$197,885.20\end{aligned}$$

What these results tell us is that purchasing machine A is equivalent to receiving $193,755.66 every year, whereas purchasing machine B is equivalent to receiving $197,885.20 each year. Clearly, machine B is preferable.

Year	Machine A Cash Flows ($)	Machine B Cash Flows ($)
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 8: Machine A and B Free Cash Flows Over Least Common Multiple Lifetime

Copyright © 2020 Michael R. Roberts

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 8, 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 equivalent cash flows over a 12-year period, as the following calculations show.

$$\text{Machine A NPV} = \frac{193,755.66}{0.10}\left(1-(1+0.10)^{-12}\right) = \$1,320,191.33$$
$$\text{Machine B NPV} = \frac{197,885.20}{0.10}\left(1-(1+0.10)^{-12}\right) = \$1,348,328.74$$

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.

As a follow up on Roarke's problem, consider what would happen if they had an existing machine in place that would serve for another two years and generate $500,000 one year from today, $200,000 two years from today. They can postpone investing in either machine A or B and keep using this old machine. When should they replace the existing machine? This question is answered by comparing the cash flows to the annuity equivalent cash flows.

Next year the existing machine will earn $500,000, whereas machines A and B offer annuity equivalents less than $200,000. Remember, we don't have to spend any money on our existing machine, that's a sunk cost and as such irrelevant for our decision. Likewise, two years from today, the old machine will generate $200,000 in cash flow, again greater than the annuity equivalent of the new machines. Ultimately, we want to use up that machine until it can no longer generate annual cash flows that are greater than the new machines, which takes place after two years. (This assumes that the costs and benefits of the replacement machines don't change as we delay replacement.)

6.4 Application: Customer Lifetime Value (CLV)

In the late 1980s, marketing executives realized that they could apply discounted cash flow analysis to customers. In essence, each customer is like a small project generating a sequence

Copyright © 2020 Michael R. Roberts

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 larger project (e.g., Dell Tablet) or even an entire company.

While there is debate over what this aggregation 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.

6.4.1 The Basic Model

Perhaps the most basic version of a customer lifetime model is the following.

$$CLV = \text{Duration} \times \text{Recurring revenue} \times \text{Operating margin} \tag{6.2}$$

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) company we'll call Cuesta focused on B2B (business to business) 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, which hosts Cuesta's software and its clients' data. 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

Copyright © 2020 Michael R. Roberts

marketing, advertising, sales expenses, etc.

$$CLV = -CAC + \text{Duration} \times \text{Recurring revenue} \times \text{Operating margin} \tag{6.3}$$

For our example, let's assume the CAC is \$5,000 per customer, implying a customer lifetime value net of customer acquisition costs of $17,280 - 5,000 = \$12,280$.[1]

Some common key performance indicators for CLV analysis include.

1. **LTV-to-CAC Ratio**

$$\text{LTV-to-CAC} = \frac{17,280}{5,000} = 3.46\text{x}$$

This measure tells us that a customer generates \$3.46 for each dollar we have to spend acquiring 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 periods (months in this example) to recover the money spent acquiring the customer (CAC) from the profits we receive from them.

3. **Churn Rate**

$$\textbf{Churn} = \frac{1}{\text{Expected customer lifetime}} = \frac{1}{36} = 0.0278$$

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.

6.4.2 DCF in Disguise

Though it may look and sound different, customer lifetime value is just discounted cash flow analysis by another name. Let's write the CLV expression (equation 6.3) in a slightly

[1]Whether the CAC is OpEx or CapEx depends on the nature of the investment and accounting rules. This distinction could be important for tax reasons.

Copyright © 2020 Michael R. Roberts

different manner so we can more easily see the connections.

$$\begin{aligned} CLV &= \underbrace{\text{Customer Acquisition Costs}_0}_{CashFlow_0} \\ &+ \underbrace{(\text{Sales}_1 - \text{Operating expenses}_1)}_{CashFlow_1} + \underbrace{(\text{Sales}_2 - \text{Operating expenses}_2)}_{CashFlow_2} + ... \\ &+ \underbrace{(\text{Sales}_T - \text{Operating expenses}_T)}_{CashFlow_T} \end{aligned}$$

In other words, equation (6.3) is just 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 real question is whether these assumptions make any 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. Taxes could also be zero if Cuesta is unprofitable.

However, no investment in net working capital is a stretch unless the customer pays cash every month. This seems unlikely in our Cuesta example or even a B2C (business to customer) setting in which customers pay by credit card. Perhaps most troubling is the lack of discounting. Even if the cash flows are risk-free because the relationship is governed by a strict contract, the discount rate should not be zero. It should be the risk-free rate!

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 CLV analysis.

Copyright © 2020 Michael R. Roberts

6.5 Key Ideas

We now have a largely complete view of corporate capital budgeting, which is really no different from personal capital budgeting, i.e., how we should allocate our own money.

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.
- 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 Problems

For all problems requiring calculation, it is strongly recommended - in many cases required - that a spreadsheet or other computing program be used.

Copyright © 2020 Michael R. Roberts

6.1 (*IRR sign rule*) The following table presents cash flows for six projects.

	Year				
Project	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.

	Year			
Project	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.

Copyright © 2020 Michael R. Roberts

a. What are the IRR and NPV of the project if Exxon-Mobil chooses *not* to finance the deal through the Brazilian central bank?

b. What are the IRR and NPV of the project if Exxon-Mobil chooses to finance the deal through the Brazilian central bank?

c. 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%

d. 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?

e. 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.

	Year					
Project	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.

a. What are the IRR, NPV, and profitability index (PI) for each project?

b. 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.

Copyright © 2020 Michael R. Roberts

	Year			
Projects	0	1	2	3
Delaware	(120.00)	90.00	90.00	90.00
Montana	(40.00)	30.00	30.00	30.00
Oregon	(50.00)	20.00	20.00	20.00
Vermont	(30.00)	10.00	10.00	10.00

The cost of capital for each boat is 15%.

Using this information, answer the following questions.

a. 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.

Copyright © 2020 Michael R. Roberts

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

Copyright © 2020 Michael R. Roberts

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.

a. Compute the net present values for the three replacement times.
b. Which replacement time offers the largest NPV? Is this the optimal replacement time?
c. 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.

Copyright © 2020 Michael R. Roberts

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%.

Copyright © 2020 Michael R. Roberts

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.

a. What is the annual operating cost for each light bulb?

b. What is the present value of the total costs - purchase plus operating - of one lifetime for each bulb?

c. Using your answer from the previous question, what are the periodic equivalent cash flows for each bulb?

Copyright © 2020 Michael R. Roberts

d. 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)?

e. 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?

f. (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.

a. What are the hourly and daily operating costs for each light bulb?

b. How many days and years will each bulb last?

c. What is the present value of the total costs of one lifetime for each lightbulb? (Hint: Assume a period is one day.)

d. Using your answer from the previous question, what are the periodic equivalent cash flows for each bulb?

e. 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)?

Copyright © 2020 Michael R. Roberts

f. 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?

g. (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?

h. 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	450	3	2,795	9.0	18
3	100	2	1,400	15.0	12
4	1,000	4	4,700	3.0	24
5	300	2	1,600	12.0	12

Using this information, answer the following questions.

(a) 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.

(b) Derive closed-form expressions (i.e., equation) for the payback period and the discounted payback periods and use your results to estimate these measures.

Copyright © 2020 Michael R. Roberts

(c) Compute the profitability index of one tenant cycle for each apartment using prep cost as the denominator. Assuming management only has $36,500 of cash on hand for prep costs, which apartments should it ready today according to your profitability index calculations?

(d) Compute the profitability index of one tenant cycle for each apartment assuming using crew size as the denominator. Rank the apartments based on this measure. Assuming management only has seven 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. (Careful...the units have different costs of capital.) Revisit your answers to the two previous questions under the assumptions that the apartments will be rented in perpetuity. Would you prepare the same apartments when facing either a monetary or head-count constraint?

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.

Copyright © 2020 Michael R. Roberts

- 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

Copyright © 2020 Michael R. Roberts

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

on bankrate.com, Luna has assembled the following list of mortgage products from which to choose.

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

Copyright © 2020 Michael R. Roberts

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?

Copyright © 2020 Michael R. Roberts

Chapter 7

Investing: Bonds

Fundamental value relation

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

This chapter

- 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 each period as investors, and realized returns, what we actually earn each period,
- considers the implications of taxes on investor returns,
- identifies and quantifies the risks associate 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?

Copyright © 2020 Michael R. Roberts

- How can we hedge, 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?

Bonds are a subset of loans often referred to as *fixed income instruments*. In fact, we've already discussed several fixed income instruments including fixed-rate mortgages, corporate bonds, and treasury securities. It should come as no surprise that we'll use our fundamental valuation relation to value these instruments. What will be new is that we'll use different discount rates, r, to discount different cash flows. These rates will come from the term structure of interest rates that we discussed in chapter 3. Also new is that we'll pay close attention to the risks of investing in fixed income instruments, notably interest rate, default, and inflation risk.

7.1 Zero Coupon Bonds

We'll start off by examining the most basic bond, a **zero coupon bonds** or more simply a **zero**. A zero is a loan in which repayment is made in one lump sum at the end or **maturity** of the bond. How much is repaid at maturity is the **face** or **par** value of the bond. **Treasury bills**, or **T-bills** as they're often called, are zero coupon bonds with maturities less than one year. There are five different T-bills issued by the U.S. treasury and differentiated by their maturities: 4-weeks, 8-weeks, 13-weeks, 26-weeks, and 52-weeks. Corporations issue zero coupon bonds, but only rarely for reasons that will be apparent later.

7.1.1 Valuation

Consider a 12-month zero coupon bond with a par value of $1,000. If the APR is 1%, what should the price of the bond be today if interest rates are compounded annually? Let's denote the price at time t by P_t, so the price today is represented by P_0. As always, start with a timeline.

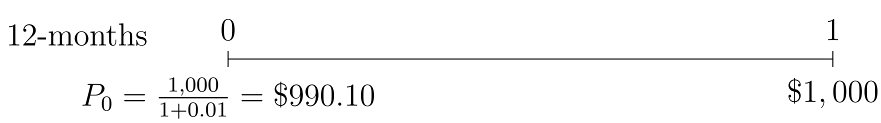

Figure 7.1: Pricing a One-Year Zero Coupon Bond

Copyright © 2020 Michael R. Roberts

Figure 7.2: Pricing a One-Year Zero Coupon Bond

If we had a 10-year zero coupon bond with a $1,000,000 face value and a 4% APR compounded semi-annually, the timeline and price are show in figure 7.2. From the last chapter, we could also compute the bond price using the effective annual rate, $r = (1 + 0.04/2)^2 - 1 = 4.04\%$.

$$P_0 = \frac{1,000,000}{(1 + 0.0404)^{10}} = \$672,971.3331$$

Note how the price of a zero coupon bond will always be less than the face or par value of the bond as long as interest rates are positive.[1] Zero coupon bonds are referred to as **discount bonds** to reflect the typical relation between the price and face value of the bond.

A common convention in fixed income markets which applies to all bonds, not just zero coupon bonds, is to quote bond prices per $100 of face value. Thus, the price of our one-year bond would be quoted as $99.01, the price of our 10-year bond as $67.30.

7.1.2 Yield to Maturity

Now let's find the APR for a one-year zero coupon bond that is priced at $98.45 per $100 of face value and experiences annual compounding. We'll denote the face value of a bond by F. Using our fundamental valuation relation

$$P_0 = \frac{F}{(1 + r)^T} \implies r = APR = \left(\frac{F}{P_0}\right)^{1/T} - 1.$$

Plugging numbers in produces

$$APR = \left(\frac{100}{98.45}\right)^{1/1} - 1 = 0.0157.$$

The APR of a bond is also called the **yield to maturity** or just **yield**. The yield on a bond is the one interest rate such that when we discount all of the promised cash flows,

[1] Nominal interest can and have gone negative in many instances, though it is fairly rare. Interest rates on government debt in Switzerland, Denmark, Sweden, and the Euro area were all negative at some point in 2016. The German government made billions in 2021 because of negative interest rates on its debt. Negative rates can be viewed as the price investors are willing to pay governments, and in some cases companies, to hold their cash.

Copyright © 2020 Michael R. Roberts

we obtain the bond price. This should sound familiar. The yield on the bond is the same thing as the internal rate of return on the bond. Move the price to the right hand side of the equation and we get an expression that equates the NPV of a bond investment to zero. In our zero coupon setting, there is only one cash flow. Below we'll see bonds with multiple cash flows.

More generally, the yield, y, on a zero coupon bond can be computed as

$$P_0 = \frac{F}{(1+y/k)^{k\times T}} \implies y = \left(\left(\frac{F}{P_0}\right)^{1/(k\times T)} - 1\right) \times k. \tag{7.1}$$

where k is the number of compounding periods per year and T is the number of years to maturity. For example, the yield on a 3-year zero coupon bond with \$1,000 face value, semi-annual compounding, and a price of \$924 is

$$\left(\left(\frac{1,000}{924}\right)^{1/(2\times 3)} - 1\right) \times 2 = 0.0265.$$

Understand that bond yields come from bond prices, not the other way around. Supply and demand in bond markets determine bond prices, which in turn determine bond yields. Also recognize that bond prices and yields contain the same information. In other words, if we know the yield on a bond, we can derive its price and vice versa.

7.1.3 Yield Curve

If we plot the yields of bonds with different maturities against their maturities, we create a **yield curve**. An example, is presented in figure 7.3, which presents the Treasury yield curve from January 3, 2022. A yield curve is just a visualization of the term structure of interest rates that was discussed in Chapter 3. Thus, for each term structure of interest rates - Treasuries, mortgages, corporate bonds, etc. - there is a corresponding yield curve.

We can use a yield curve to identify discount rates for cash flows at different points in time. For example, imagine today is January 3, 2022 and we own a 5-year **Treasury STRIP** - a zero coupon bond constructed from the coupon of a longer term Treasury - whose face value is \$1,000. To price this bond, we need a discount rate for the cash flow we'll receive five years from today. Looking at the yield curve in figure 7.3, we see the 5-year yield is 1.37%.

If the yield is semi-annually compounded, then the price of our bond is

$$P_0 = \frac{1,000}{(1+0.0137/2)^{10}} = \$934.01$$

Copyright © 2020 Michael R. Roberts

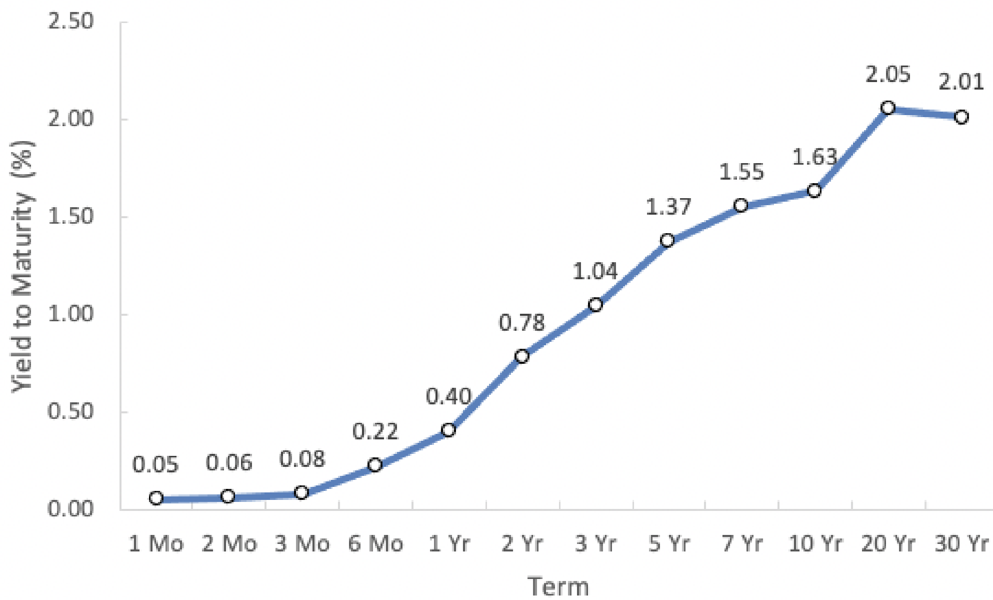

Figure 7.3: Treasury Yield Curve, January 3, 2022 (Source: US Department of Treasury)

Had our bond maturity been 20 years, its yield would have been 2.05% and it's price, again assuming semi-annual compounding, would be

$$P_0 = \frac{1,000}{(1 + 0.0205/2)^{40}} = \$665.04$$

7.1.4 Price Dynamics and Returns

What happens to the price of a bond after buying it depends largely on what happens to interest rates, inflation, and **default risk**, or the risk that the borrower fails to repay the loan. To avoid complicating things at this point, let's ignore the latter two risks and examine Treasury STRIP with five years to maturity.[2] Let's also assume that interest is compounded annually, again to avoid unnecessary complications.

The price of our bond today is

$$P_0 = \frac{100}{(1 + 0.0137)^5} = 93.42.$$

The discount rate, 1.37%, is the 5-year yield found in figure 7.3.

To compute the price one year later, we need to know what the four-year yield will be one year from today. We need the four-year yield because that is how much time remains

[2]Strictly speaking, there is some default risk even with Treasury securities as the U.S. government could decide to not pay its debt. Historically, thee probability of this event was extremely small.

Copyright © 2020 Michael R. Roberts

until we receive the face value payment. The problem is that we don't know what the four-year yield will be one year from today because we don't know what interest rates will be in the future. So, in practice, we have to make an educated guess or build models to forecast interest rates - exactly what professional bond investors do.

For now, let's pretend that we do know what the four-year yield, one year from today is and that equals 1.205%. In this case, the price of our 5-year Treasury STRIP one year from today will be

$$P_1 = \frac{100}{(1+0.01205)^4} = \$95.32.$$

If we choose to sell the bond one year from today at this price, then our one-year realized return will be

$$r_{0,1} = \frac{P_1}{P_0} - 1 = \frac{95.32}{93.42} - 1 = 0.0203,$$

where $r_{0,1}$ denotes the **net return** or simply **return** from period 0 to period 1. In words, the return is just the ratio of the money we make (P_1 if we sell the bond) divided by the money we spent (P_0 when we bought the bond). Returns tell us how much money we make when positive (lose when negative) after accounting for how much we invested (the minus one part).

The 2.03% return in this example implies that a one dollar investment in the bond generates an additional \$0.0203. When people talk about returns they are almost always talking about net returns, numbers like 10% or -4%. **Gross returns** don't subtract "1" and are numbers like 1.10 or 0.96. For example, a gross return of 1.10 tells us for each dollar we invest, we get \$1.10 back, which includes our original investment.

A couple of comments regarding the return calculation

- Useful pneumonics to remember how to compute **net returns** is

$$\text{Return} = \frac{\text{Money made}}{\text{Money spent}} - 1 = \frac{\text{New Price}}{\text{Old Price}} - 1 \tag{7.2}$$

  Gross returns don't subtract one.

- Returns are earned over a period of time or horizon (e.g., year, month, etc). We'll explicitly indicate this period using subscripts. For example, $r_{t-1,t}$ refers to the return from period $t-1$ to t.

- The one-year return, 2.03%, is higher than the original yield on the bond when we bought it, 1.37%. The reason for this increase is that the interest rate "fell" from 1.37% to 1.205%. Consequently, the bond price increased by more than expected from the passage of time.

Copyright © 2020 Michael R. Roberts

This last point deserves a closer look. Let's recompute the price of the bond with four years to maturity assuming now that the 4-year yield one year from today is the same as the 5-year yield today when we bought the bond, 1.37%.

$$Price_1 = \frac{100}{(1 + 0.0137)^4} = \$94.70$$

Our one-year return is now

$$Return_{0,1} = \frac{94.70 - 93.42}{93.42} = 0.0137,$$

identical to the bond yield when we bought it. Because interest rates didn't change, the return on our bond is exactly the same as its current yield.

The lesson here is that **bond prices increase when interest rates decrease and vice versa**. Table 1 illustrates price dynamics and returns over the life of our 5-year Treasury STRIP for some randomly chosen prices. To be clear, the prices for years one through four are *future* bond prices that we are pretending to know. From these prices, we can compute the corresponding bond yields.

	Period (Year)					
	0	1	2	3	4	5
Yield (%)	1.37	1.21	1.80	2.10	0.40	
Price ($)	93.42	95.32	94.79	95.93	99.60	100.00
	Annualized Returns (%)					
1-year		2.03	-0.56	1.20	3.83	0.40
2-year			0.73	0.32	2.51	2.10
3-year				0.89	1.47	1.80
4-year					1.61	1.20
5-year						1.37

Table 1: Zero Coupon Bond Price and Annualized Return Dynamics

The Yield row in the table shows how interest rates evolve over time. Today, period 0, the 5-year yield is 1.37%. Next year, period 1, the 4-year yield is 1.21%. And so on. The

Copyright © 2020 Michael R. Roberts

following equalities show the relation between the bond prices and yields.

$$P_0 = \frac{100}{(1+0.0137)^5} = \$93.42$$
$$P_1 = \frac{100}{(1+0.0121)^4} = \$95.32$$
$$P_2 = \frac{100}{(1+0.0180)^3} = \$94.79$$
$$P_3 = \frac{100}{(1+0.0210)^2} = \$95.93$$
$$P_4 = \frac{100}{(1+0.0040)^1} = \$99.60$$

At maturity, year 5, the bond price equals the face value because there is no more waiting to receive the money.

The bottom section of the table presents annualized returns over different horizons. We saw how to compute the 1-year return from period 0 to 1 above. The return from period 1 to 2 is:

$$r_{1,2} = \frac{94.79 - 95.32}{95.32} = -0.0056.$$

This negative return suggests an investor that bought this bond one year from today and then sold it a year later lost money. Indeed, the price of the bond is lower in year 2 than year 1. This decline is a result of a large increase in interest rates between periods 1 and 2, as seen in the change in bond yields from 1.21% to 1.80%.

Multi-year returns are computed in the same manner as one year returns (equation 7.2). For example, the 3-year return from today to three years later is

$$r_{0,3} = \frac{95.93 - 93.42}{93.42} = 0.0268.$$

It's important to emphasize that this is the return over a 3-year horizon. As such, it shouldn't be compared with one-year returns. To ease comparisons across different horizons, the table presents **annualized returns** defined as follows.

$$\text{Annualized return} = (1 + r_{t,T})^{(1/(T-t))} - 1 \tag{7.3}$$

The annualized return from today to three years later is

$$\left(\frac{95.93}{93.42}\right)^{1/3} - 1 = 0.0089,$$

as seen in the table.

Copyright © 2020 Michael R. Roberts

To conclude, bond prices change because interest rates change. Consequently, buying and selling bonds, even bonds with guaranteed payments, comes with risk if we choose to sell the bond at any point prior to maturity. For risk-free bonds like Treasuries, if we buy and hold the bond until maturity, that interest rate risk is eliminated. The bottom row of table 1 illustrates this elimination. The annualized 5-year yield on the bond is 1.37%, the yield we computed today.

7.1.5 Expected Returns vs. Realized Returns

The returns computed in table 1 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 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, P_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 $\mathbb{E}(P_1)$. This expression reads "the expected value of P_1 (i.e., the price one year from today)." The **expected return** of our bond investment over the next year is

$$\mathbb{E}(r_{0,1}) = \frac{\mathbb{E}(P_1) - P_0}{P_0}.$$

Because the expectation notation, $\mathbb{E}$, can appear daunting, I'll suppress it when it's clear that we're working with expected values.

The only difference from our realized return calculation is the replacement of the P_1 with $\mathbb{E}(P_1)$. We'll talk more later about how to compute expected values. 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 take an educated guess to estimate those values.

Of course, this is nothing new. Our earlier discussions about whether or not to go to college (Chapter 2) and corporate decision making (Chapter 5) both implicitly used 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 decisions. We have to estimate

Copyright © 2020 Michael R. Roberts

all of these data, and we do so using expected values. All we've done in this section is be explicit in recognizing that we have to estimate unknown future quantities, and we have to differentiate between our guesses of future quantities (expectations) and what those future quantities eventually turn out to be (realizations).

7.1.6 Taxes on Bond Returns

Earnings on bond investments are taxed in the U.S. Specifically, 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. Reconsider the 10-year, zero-coupon bond with a face value of $1 million and 4% semi-annual compounded APR discussed above. The bond value today is $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 were 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 our 10-year 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

Copyright © 2020 Michael R. Roberts

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.

7.1.7 Market Efficiency and Timing

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 market price.[3] An implication of market efficiency is that 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 our investment, because the market knows everything you 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." The argument that interest rates are low is suggestive that they will go up again in the future. So, we'd better buy a house while we can get cheap financing, i.e., before interest rates rise.

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 priced into financial instruments. That is to say, 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 not probably not getting a special deal on mortgages because the interest rate appears low relative to historical standards.

In fact, as figure 7.4 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 mortgages and below 2% on some adjustable rate mortgages.

[3]Technically, 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.

Copyright © 2020 Michael R. Roberts

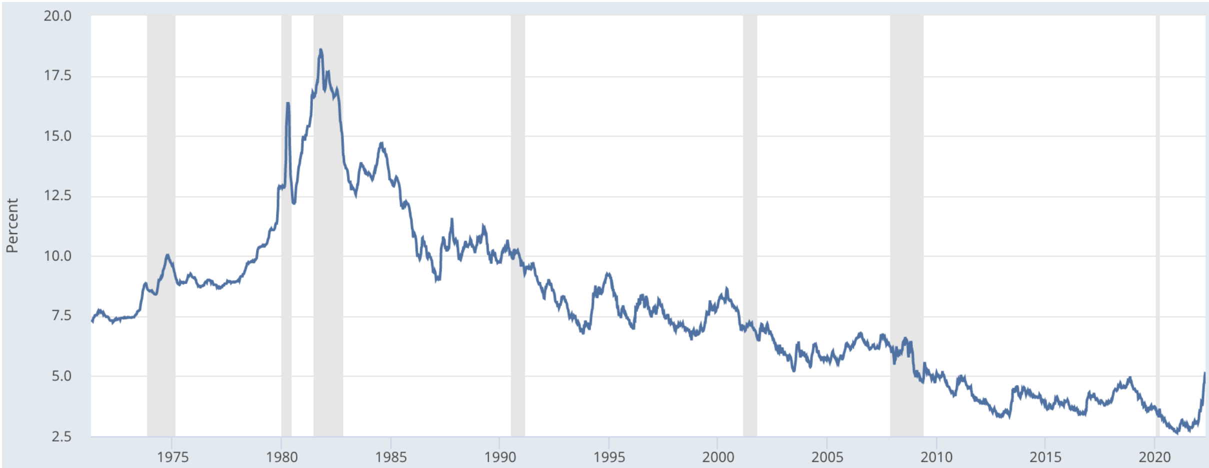

Figure 7.4: 30-Year Fixed Rate Mortgage Average in the United States, 2003-2021 (Source: St. Louis Federal Reserve - Freddie Mac)

Now consider what happened in December of 2008 after the onset of the 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 7.5 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.

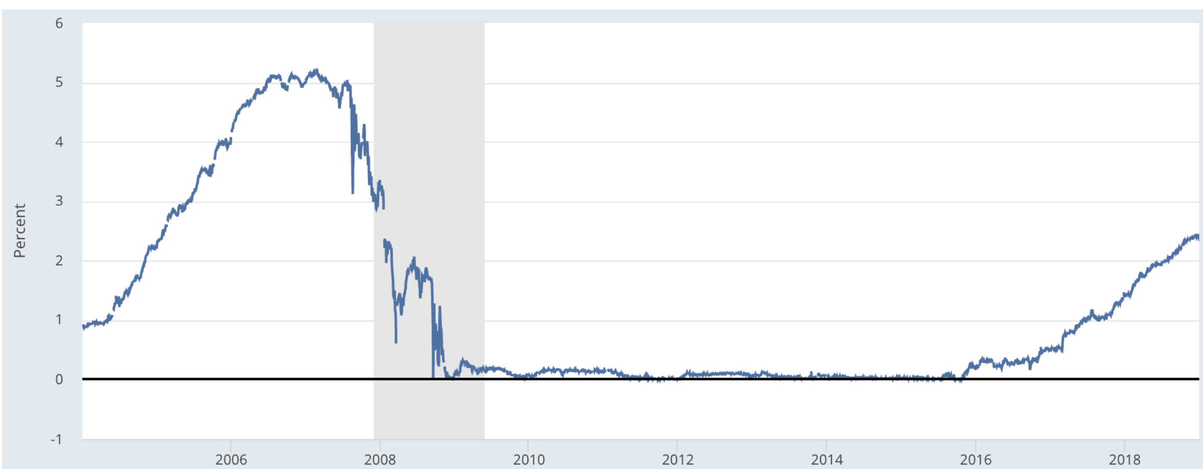

Figure 7.5: Market Yield on U.S. Treasury Securities at 3-Month Constant Maturity, 2004-2018 (Source: St. Louis Federal Reserve - Board of Governors of the Federal Reserve System (US))

When we look at market interest rates and think, “boy, these rates sure seem low (or high). 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

Copyright © 2020 Michael R. Roberts

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.

It is important not to 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? If nothing, then all these expectations about what may or may not happen in the future are already represented in the current price, which makes it difficult to make abnormal profits in financial markets, where abnormal means beyond what we should get for the risk we take.

7.2 Coupon Bonds

Unlike a zero, a **coupon bond** makes periodic interest payments over the life of the loan, and a final payment of principal (i.e., par or face value) at maturity. Some examples of coupon bonds include:

- **Treasury notes** or **T-notes** are medium-term bonds paying a semi-annual coupon. Treasury note maturities consist of 2, 3, 5, 7, and 10 years.
- **Treasury bonds** or **T-bonds** are long-term bonds paying a semi-annual coupon with maturities of 20 and 30 years.
- **Corporate bonds** come in a variety of maturities and most often offer a semi-annual coupon. However, there are examples of zero coupon corporate bonds, as well as floating rate corporate bonds.

Let's explore the valuation of these instruments.

Copyright © 2020 Michael R. Roberts

7.2.1 Valuation

Consider a 2 year T-note with a \$100 face value and a coupon rate of 5%. To compute the semi-annual coupon payments, we multiply the face value of the bond by one half times the **coupon rate**. The coupon rate is halved because it is expressed as an APR. The periodic coupons are the coupon rate divided by the coupon frequency. The coupon payments for our 2-year T-note are $100 \times 0.05/2 = \$2.50$. The last payment occurring at maturity consists of both a coupon payment and the face value. The timeline of cash flows is show in figure 7.6

6 months	0	1	2	3	4
Cash flows		2.50	2.50	2.50	\$102.50

Figure 7.6: Two-Year Treasury Note Cash Flows

Valuation of this bond is like valuation of any other asset. We use our fundamental valuation relation to discount and sum the future cash flows. To do so, we need a discount rate(s). Let's use the Treasury yield curve to get the discount rates. Referring back to figure 7.3, we see that the yields for 6-month, 12-month, and 24-month Treasuries are 0.22%, 0.40%, and 0.78%, respectively. There is no yield for the 18-month payment in the figure so let's take the midpoint between the 12- and 24-month yields, $(0.78 + 0.40)/2 = 0.59$.

The price of our bond should be

$$Price_0 = \frac{2.50}{(1 + 0.0022/2)} + \frac{2.50}{(1 + 0.0040/2)^2} + \frac{2.50}{(1 + 0.0059/2)^3} + \frac{102.50}{(1 + 0.0078/2)^4} = \$108.38$$

Note that each yield, expressed as an APR, needs to be divided by two to get the periodic interest rate. This is just convention.

As we did with the zero coupon bonds, we can use the bond price to find the bond's yield to maturity. The yield to maturity for this 2-year T-note can be found by solving

$$108.38 = \frac{2.50}{(1 + y/2)} + \frac{2.50}{(1 + y/2)^2} + \frac{2.50}{(1 + y/2)^3} + \frac{102.50}{(1 + y/2)^4}$$

for y. To do so, we need a computer - Excel's IRR or Goal Seek functions will work. The yield on this bond is 0.7688%.

A little bond lingo: When the bond price is greater than the par value, the bond is said to be **priced at a premium** and the bond yield will be less than the coupon rate. In this example, the bond is priced at a premium (\$108.38 compared to \$100 par value), and its yield is less than the coupon rate (0.7688% compared to 5%). When the bond price is less

Copyright © 2020 Michael R. Roberts

than the par value, the bond is said to be **priced at a discount** and the bond yield will be greater than the coupon rate. When the bond price equals the par value, the bond is said to be **priced at par** and the bond yield will equal the coupon rate.

7.2.2 Returns

The realized returns on a coupon bond are computed in a similar manner to those of a zero coupon bond. One difference is that we have to account for the coupon payments. So, the return equation we saw earlier (equation 7.2) needs to be tweaked like so.

$$r_{t-1,t} = \frac{\text{Money made}}{\text{Money spent}} - 1 = \frac{P_t + Coupon_t}{P_{t-1}} - 1 \tag{7.4}$$

Table 2 examines the future prices and realized returns of our two-year T-note assuming interest rates don't change - i.e., the yield curve is constant over time. Bond prices are computed using our fundamental valuation relation. For

$$\begin{aligned}
P_1 &= \$105.68 = \frac{2.50}{(1+0.0022/2)} + \frac{2.50}{(1+0.0040/2)^2} + \frac{102.50}{(1+0.0059/2)^3} \\
P_2 &= \$104.18 = \frac{2.50}{(1+0.0022/2)} + \frac{102.50}{(1+0.0040/2)^2} \\
P_3 &= \$104.18 = \frac{102.50}{(1+0.0022/2)}
\end{aligned}$$

For one period returns, we use equation 7.4. More generally, we compute returns as internal rates of return on the bond to deal with multiple cash flows occurring over time (i.e., the coupon payments). For example, if we buy the bond today and sell it six months from today after receiving the coupon payment, our six-month return is

$$108.38 = \frac{2.50 + 105.68}{(1 + r_{0,1})} \implies r_{0,1} = -0.0018$$

The annualized return is $(1 - 0.0018)^2 - 1 = -0.0037$. If instead we buy the bond one year (two periods) from today and sell it six months later, our six-month return is

$$104.18 = \frac{2.50 + 102.27}{(1 + r_{0,1})} \implies r_{2,3} = 0.0057,$$

or annualized, $(1 + 0.0057)^2 - 1 = 0.0115$.

Copyright © 2020 Michael R. Roberts

	Period (6-months)				
	0	1	2	3	4
Yield curve (%)		0.22	0.40	0.59	0.78
Treasury note cash flows		2.50	2.50	2.50	102.50
Prices	108.38	105.68	104.18	102.27	
6-months					
Cash flows	(108.38)	108.18			
Return (%)	(0.37)				
Cash flows		(105.68)	106.68		
Return (%)		1.90			
Cash flows			(104.18)	104.77	
Return (%)			1.15		
Cash flows				(102.27)	102.50
Return (%)				0.44	
1-year					
Cash flows	(108.38)	2.50	106.68		
Return (%)	0.74				
Cash flows		(105.68)	2.50	104.77	
Return (%)		1.53			
Cash flows			(104.18)	2.50	102.50
Return (%)			0.80		
18-months					
Cash flows	(108.38)	2.50	2.50	104.77	
Return (%)	0.88				
Cash flows		(105.68)	2.50	2.50	102.50
Return (%)		1.17			

Table 2: Coupon Bond Price and Annualized Return Dynamics - Fixed Yield Curve

Copyright © 2020 Michael R. Roberts

Longer-term returns are computed similarly. For example, if we purchase the bond six months from today and sell it one year later, our periodic return is (using Excel)

$$105.68 = \frac{2.50}{(1+r_{1,3})} + \frac{2.50 + 102.27}{(1+r_{1,3})^2} \implies r_{1,3} = 0.0076,$$

or annualized, $(1 + 0.0076)^2 - 1 = 0.0153$.

An important reminder is that the calculation of the internal rate of return assumes that any cash flows arriving before the end of the investment are reinvested at the internal rate of return. Using our previous example of purchasing a bond six months from today and selling it one year later, the \$2.50 coupon we receive six months after purchasing the bond needs to be reinvested. The internal rate of return assumes it is reinvested at 0.76% per six months or 1.53% per annum. If this is *not* the case, then the internal rate of return is not giving us a precisely accurate picture of our investment return.

7.2.3 Taxes

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 above. To illustrate, consider the following example.

Imagine we purchase the two-year treasury note from table 2 for \$108.38 today, and then sell it for \$102.27 18 months later. Also assume that we purchase the bond on December 31, 2021 and 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 7.7.

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 7.7: 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 $108.38 - 102.27 = \$6.11$,

Copyright © 2020 Michael R. Roberts

which we can deduct from our taxable income.[4] 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).

Because taxes can significantly alter our investment returns, alternatives that avoid taxes, partially or fully, are popular among investors in high tax brackets. For example, income from Treasury securities are not taxed at state and local levels, but they are at the federal level which tends to represent the lion's share of the tax burden. **Municipal bonds** or **munis** are tax-free at *all* levels - local, state, and federal - provided we reside in the state in which they are issued. For example, a Pennsylvania resident that purchases a Pennsylvania municipal bond collects interest tax-free. However, a Pennsylvania resident that purchases a California municipal bond may be required to pay Pennsylvania taxes. For a Texas resident that faces no state taxes, it doesn't matter which municipality in which the bond was issued, they will face no taxes on any interest.

Because of their preferential tax treatment, munis tend to trade at lower yields than otherwise similar Treasury securities. This relation can be seen in figure 7.8. Municipal bond yields are lower Treasury yields for each bond maturity. This difference is due largely to tax differences, though other considerations including credit and liquidity risk are responsible for some of the differences. Specifically, municipal bonds can and do default (i.e., fail to repay their owners) at a greater frequency than the federal government.[5] Likewise, municipal bonds are less liquid - i.e., less frequently traded and more costly to trade - than their Treasury counterparts. Both of these risks offset the tax benefits that munis provide.

7.2.4 Replication

Another way to value our T-note is as the price of a portfolio of assets whose cash flows exactly match those of our T-note. (A **portfolio** is a collection of assets.) One way to view our T-note is as a combination of (i) an annuity and (ii) a zero coupon bond. The coupons correspond to an annuity, the face value to a zero coupon bond. If we know the price of the annuity and the zero coupon bond, we can add them to get the price of our T-note. Table 3 illustrates the logic.

[4] As 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.

[5] Some notable recent municipal bond defaults occurred in California (Orange County in 1994, Vallejo in 2008, Stockton in 2013), Alabama (Jefferson County in 2008), and Pennsylvania (Harrisburg in 2011). During the Great Depression, over 4,700 municipalities defaulted during the 1930s.

Copyright © 2020 Michael R. Roberts

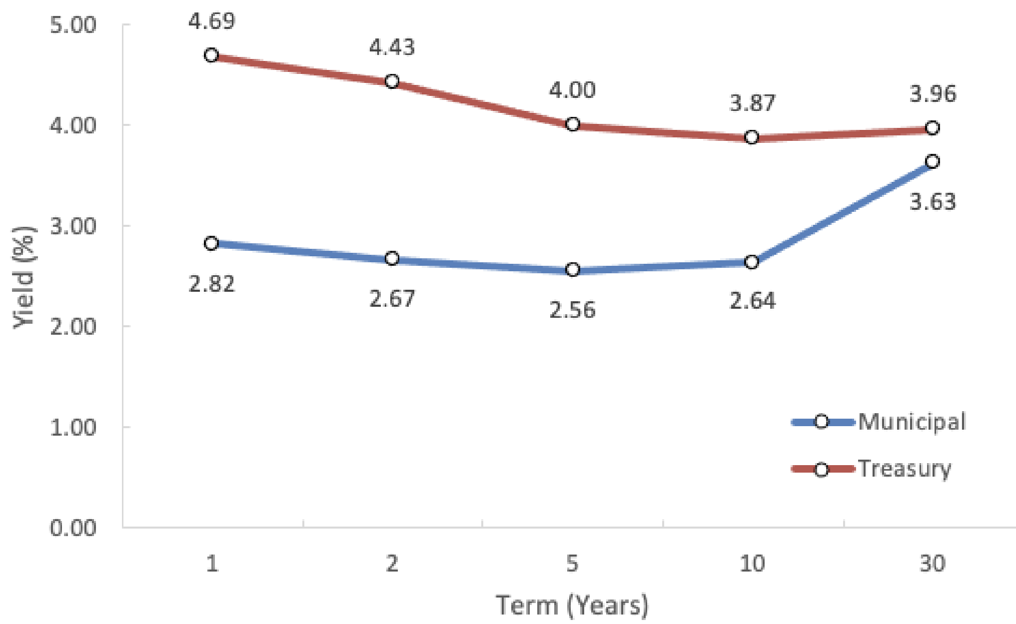

Figure 7.8: Municipal and Treasury Bond Yield Curves, December 31, 2022 (Source: Bloomberg)

	Periods (Six Months)				
	0 ($Price_0$)	1	2	3	4
T-note	?	2.50	2.50	2.50	102.50
Annuity	9.93	2.50	2.50	2.50	2.50
Zero	98.45				100
Replicating Portfolio	108.38	2.50	2.50	2.50	102.50

Table 3: Treasury Note Replicating Portfolio (Annuity & Zero)

The first row of the table identifies the periods, 0 through 4, representing semi-annual increments. The second row shows the T-note we want to price and its future cash flows. The third and fourth rows show the prices and future cash flows of two assets. The first is a two-year, semi-annual annuity paying $2.50. The price of this asset can be found using the term structure in figure 7.3

$$Price_0^{Annuity} = \frac{2.50}{(1+0.0022/2)} + \frac{2.50}{(1+0.0040/2)^2} + \frac{2.50}{(1+0.0059/2)^3} + \frac{2.50}{(1+0.0078/2)^4} = \$9.9267$$

The second asset is a four-year zero coupon bond with a face value of $100. The price of this asset can be found by discounting the face value with the two-year yield, 0.78%

Copyright © 2020 Michael R. Roberts

$$Price_0^{Zero} = \frac{100}{(1+0.0078/2)^4} = \$98.4551$$

The sum of the annuity and zero cash flows exactly match the cash flows of the 2-year T-note. As such, the portfolio containing the annuity and zero coupon bond is called a **replicating portfolio**. A replicating portfolio generates future cash flows that exactly match - in terms of timing and magnitude - those of another asset, in this case, the T-note. This replicating portfolio is also called a **synthetic** T-note because we synthesized a T-note with other securities.

Because the future cash flows are identical under all circumstances, the prices of the T-note and the replicating portfolio should be the same. The price of the replicating portfolio is 9.9267 + 98.4551 = \$108.38, exactly what we found earlier.

Let's try this valuation approach with a different replicating portfolio, this time consisting of four zero coupon bonds. Table 4 lays out the securities and their cash flows. The zero coupon bond prices are determined using the yield curve in figure 7.3.

	Periods (Six Months)				
Security	0 ($Price_0$)	1	2	3	4
T-note	?	2.50	2.50	2.50	102.50
6-month zero	$\frac{2.50}{(1+0.0022.2)} = 2.4973$	2.50			
12-month zero	$\frac{2.50}{(1+0.0040/2)^2} = 2.4900$		2.50		
18-month zero	$\frac{2.50}{(1+0.0059/2)^3} = 2.4780$			2.50	
24-month zero	$\frac{102.50}{(1+0.0078/2)^4} = 100.9165$				102.50
Replicating Portfolio	108.38	2.50	2.50	2.50	102.50

Table 4: Treasury Note Replicating Portfolio (Zeros)

Again, the cash flows from the replicating portfolio exactly match the cash flows of our T-note. Thus, the prices should be the same.

Now, it's not always the case that there are traded assets that can precisely replicate the cash flows of another asset. But, for many financial products, particularly more complicated products (e.g., derivatives, variable rate annuities), it is easier to understand and value the product in terms of its replicating portfolio. Replicating portfolios are all also useful for executing **arbitrage strategies**, or strategies that take advantage of risk-free profit opportunities.

Copyright © 2020 Michael R. Roberts

7.3 Application: Bond Arbitrage and Short Sales

If the prices of our T-note and replicating portfolio are *not* the same, then an arbitrage opportunity exists. Let's see how this works by assuming (i) we don't have any money or assets, and (ii) our T-note is trading in the market for a price of $102.43. We'll discuss some of the practical challenges of executing an arbitrage at the end of our discussion.

At $102.43, the T-note is cheap relative to the replicating portfolio, which costs $108.38. What do we do when something is cheap? We buy it! Problem: we don't have any money. Well, we know the replicating portfolio is expensive relative to the T-note. What do we do when something is relatively 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. Remember, we assumed we don't have any money or assets.

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 $108.38, the value of the replicating portfolio. This process of 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 that I discuss below. With $108.38 in our pocket, we can now purchase the cheap T-note for $102.43, leaving us with the difference $108.38 - 102.43 = \$5.95$.

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 from the T-note we bought. We can simply hand this coupon payment over to them. Likewise, one year from today, the person from whom we borrowed the 12-month zero is expecting to receive $2.50. Again, we hand them the coupon payment from 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 $5.95 today with no money and never paid a penny. This is an arbitrage, whose strategy is illustrated in Table 5.

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 is enough to allow us to make risk-free profit with no money. It's like finding a $5.95 on the sidewalk and not having to worry about someone asking for the money back.

Copyright © 2020 Michael R. Roberts

	Periods (Six Months)				
Security	0 ($Price_0$)	1	2	3	4
Short sell 6-month zero	$\frac{2.50}{(1+0.0022.2)} = 2.4973$	-2.50			
Short sell 12-month zero	$\frac{2.50}{(1+0.0040/2)^2} = 2.4900$		-2.50		
Short sell 18-month zero	$\frac{2.50}{(1+0.0059/2)^3} = 2.4780$			-2.50	
Short sell 24-month zero	$\frac{102.50}{(1+0.0078/2)^4} = 100.9165$				-102.50
Replicating Portfolio	+108.38	-2.50	-2.50	-2.50	-102.50
Buy T-note	-102.43	+2.50	+2.50	+2.50	+102.50
Arbitrage profit	5.95	0	0	0	0

Table 5: Arbitrage Strategy

If we're really smart, we'll scale this strategy up. Rather than short-selling $108.38 of the replicating portfolio, let's short $10,838,000 of the replicating portfolio and and use the proceeds to buy $10,243,000 of the T-note. This leaves us with an arbitrage profit of $595,000 today. But, this example begs the question: Why don't scale this strategy up infinitely? We don't need any money.

The problem comes in 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 and 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 mispricings are rare and short-lived.

Short-Selling

Some comments about short-selling are in order. There are some challenges when borrowing the bonds.

1. Most large financial institutions have securities lending desks who sole purpose is to lend out the securities they and their clients own. Nonetheless, we have to hope someone is willing to lend the security to us, which is not always the case. To short sell, you have to find a willing lender.

2. Assuming we do find a willing lender, they are 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

Copyright © 2020 Michael R. Roberts

be that the interest expense on the security loans is high enough to make the arbitrage strategy no longer profitable.

3. The lender will also require us to **post collateral**, or set aside cash or other securities in an account held by a third party to ensure the lender gets repaid in case the borrower can't pay them back later. Typically, we have to post an amount equal to or greater than the amount we are short-selling. These first two points highlight that to short sell, you need money or other assets.

After we've borrowed and sold the assets, the worry isn't over.

1. What if demand for borrowing the asset we short-sold increases? The lender may decide, within the terms of the lending agreement, to increase the interest rate and erode our arbitrage profits or even reverse them to losses.

2. What if after three months the person from whom we borrowed say the 6-month zero wants to sell their bond? We have a small problem. We need to find someone else to lend us their 6-month zero so we can return the first lender's bond and maintain our arbitrage strategy. What if we can't find someone to lend us another 6-month zero? We have a big problem called a **short-squeeze** in which we have to unwind our position, i.e., buy back the six-month zero coupon bond in the market and return it to its owner. But, if this occurs after significant increases in the price of the bond, we are stuck having to buy the bond back at high prices, which could be very costly to us.

In a nutshell, short-selling is quite risky in practice. This risk was on full display in January of 2021. A number of hedge funds - institutional investors - had short sold GameStop stock in anticipation of sharp declines of its price. However, a large numbers of retail investors began buying GameStop stock in January of 2021. This buying had several implications. First, it drove up the price of GameStop stock, which imposed margin calls on investors that were short (i.e., had short-sold the stock). A margin call refers to the requirement of the investor to post more collateral to back their short position in this example. Second, the buying frenzy made it very difficult to find stock to lend. So, when short-sellers tried to cover their short positions, it was very difficult to find stock to borrow. Further, when they did find stock to borrow, the lending rate was quite high. Consequently, many hedge funds lost a great deal of money - approximately $6 billion in total - from the price runup in the stock.

Copyright © 2020 Michael R. Roberts

7.4 Interest Rate Risk

Interest rate risk refers to the sensitivity of an asset's price to changes in interest rates. Because changes in interest rates is the primary cause for changes in bond prices, understanding interest rate sensitivity is critical for bond investors. Importantly, interest rate risk is present regardless of whether or not the borrower is likely to default and, as such, affects so-called safe bonds issued by financially healthy and stable governments.

7.4.1 Intuition

To build up some intuition, consider figure 7.9 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 \$97.58 is -2.42%). The point of the picture is to highlight how different features of a bond - maturity, coupon rate, and initial yield - affect the interest rate sensitivity of the bond.

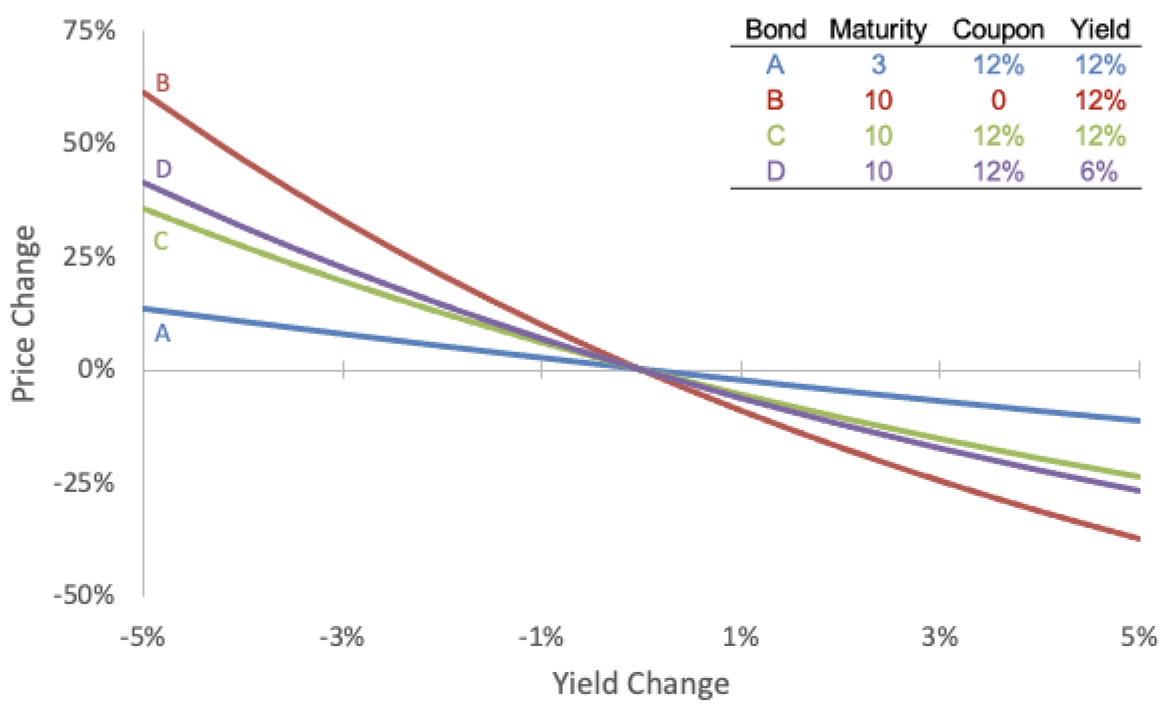

Figure 7.9: Bond Price Sensitivity to Interest Rate Changes

Several lessons can be gleaned from this figure.

Copyright © 2020 Michael R. Roberts

1. **Bond prices and yields are inversely related**. All of the curves slope down. When yields 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).

4. **Bond prices are more sensitive to decreases in interest rates than otherwise similar increases in interest rates.** Every curve is just that, a curve as opposed to a line. Each curve gets progressively less steep as we move from left to right (i.e., as interest rates increase). Mathematically, this property is referred to as **convexity**.

While this intuition is useful, we'd really like to know is how to *measure* the interest rate sensitivity of a bond. There are several ways to do so. Perhaps the most obvious is to change the interest rate a little and compute by how much the bond price changes.

7.4.2 DV01

Consider a 7-year Treasury STRIP with a \$100 face value and an annually compounded yield of 10%. The price of the STRIP today is

$$Price_0 = \frac{100}{(1+0.10)^7} = \$51.32.$$

Now let's recompute the price if the interest rate decreases by 1 basis point (bp), or 0.01%.

$$Price_0^* = \frac{100}{(1+0.0999)^7} = \$51.35$$

These computations show that 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 compute the change in bond price in response to any change in interest rates. If interest rates decrease by one point, 1%, the bond price will be

$$Price_0^* = \frac{100}{(1+0.09)^7} = \$54.70.$$

Copyright © 2020 Michael R. Roberts

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 $DV01 \times 100 = \$3.27$. 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?

$$\begin{aligned} Price_0 &= \frac{100}{(1+0.10)^{20}} = \$14.86436 \\ Price_0^* &= \frac{100}{(1+0.0999)^{20}} = \$14.89141 \\ \Longrightarrow DV01 &= \$0.02705 \end{aligned}$$

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 7.9 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 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% increases 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.[6]

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? We do this by computing the bond's yield and then changing this one number.

[6] Notice 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.

Copyright © 2020 Michael R. Roberts

Take our 2-year T-note from Section 7.2.1. As a reminder, the bond has a \$100 face value and semi-annual coupon rate of 5%. The price and yield of the bond are \$108.38 and 0.7688%. 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.

$$\begin{aligned} Price_0^* &= \frac{2.50}{(1+0.007587/2)} + \frac{2.50}{(1+0.007587/2)^2} + \frac{2.50}{(1+0.007587/2)^3} + \frac{102.50}{(1+0.007587/2)^4} \\ &= \$108.40 \end{aligned}$$

The DV01 for our coupon bond is 108.40 − 108.38 = \$0.0209. Relative to the price of the bond, a one percentage point decrease in interest rates leads to a $0.0209/108.38 \times 100 =$ 1.9284% increase in the bond price.

Let's recompute the DV01 for the same bond assuming it's coupon rate was 10% instead of 5%.

$$\begin{aligned} Price_0 &= \frac{5.00}{(1+0.007587/2)} + \frac{5.00}{(1+0.007587/2)^2} + \frac{5.00}{(1+0.007587/2)^3} + \frac{105.00}{(1+0.007587/2)^4} \\ &= \$118.3084 \\ Price_0^* &= \frac{5.00}{(1+0.007487/2)} + \frac{5.00}{(1+0.007487/2)^2} + \frac{5.00}{(1+0.007487/2)^3} + \frac{105.00}{(1+0.007487/2)^4} \\ &= \$118.3305 \\ \implies DV01 &= \$0.0221 \end{aligned}$$

Relative to the price of the bond, a one percentage point decrease in interest rates leads to a 1.8680% increase in the bond price. Again, consistent with the intuition from above, the higher the coupon rate, the less sensitive the bond is to changes in interest rates.

7.4.3 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. I've annotated the formula and written it in two different ways to help with the interpretation.

$$\begin{aligned} Duration &= \frac{1}{Price_0}\left(\overbrace{\frac{CF_1}{(1+y/k)^1}}^{PV(CF_1)} \cdot \frac{1}{k} + \overbrace{\frac{CF_2}{(1+y/k)^2}}^{PV(CF_2)} \cdot \frac{2}{k} + ... + \overbrace{\frac{CF_T}{(1+y/k)^T}}^{PV(CF_T)} \cdot \frac{T}{k} \right) \\ &= \underbrace{\frac{PV(CF_1)}{Price_0}}_{weight_1} \cdot \frac{1}{k} + \underbrace{\frac{PV(CF_2)}{Price_0}}_{weight_2} \cdot \frac{2}{k} + ... + \underbrace{\frac{PV(CF_T)}{Price_0}}_{weight_T} \cdot \frac{T}{k} \end{aligned}$$

Copyright © 2020 Michael R. Roberts

Duration tells us the weighted average of the waiting times for receiving the bonds 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 $1/k, 2/k, ...T/k$ correspond to the time, measured in years, to each cash flow. An example will be helpful.

Table 6 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.

	Periods (6-months)				
	0	1	2	3	4
Price	108.3818				
Yield	0.7688%				
Cash flows ($)		2.50	2.50	2.50	102.50
Present values, PV(CF_t) ($)		2.49	2.48	2.47	100.94
Weights, PV(CF_t) / Price		0.023	0.023	0.023	0.931
Time to cash flow (years)		0.50	1.00	1.50	2.00
Weight x Time (years)		0.01	0.02	0.03	1.86
Sum = Duration (years)	1.93				

Table 6: Duration of 2-year T-note

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

$$\frac{2.50}{(1 + 0.007688/2)^3} = \$2.47$$

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, as discussed earlier. When there are no coupon payments, i.e., a zero coupon bond, duration equals maturity.

While possibly 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

Copyright © 2020 Michael R. Roberts

measure more meaningful, we can consider **modified duration**.

$$\text{Modified duration} = \frac{\text{Duration}}{1 + y/k} \tag{7.5}$$

where y is the annual yield and k the compounding frequency. For our 2-year T-note, the modified duration is

$$\frac{1.93}{1 + 0.007688/2} = 1.92.$$

Modified duration tells us the percentage change in the price of the bond for a one percentage point (100 basis point) decrease in the interest rate.[7] For our 2-year T-note, a one percentage point decrease in interest rates leads to a 1.92% increase in the value of the bond. As long as yields are relatively small, i.e., $(1 + y/k)$ is close to one, duration and modified duration are quantitatively similar, leading practitioners to interpret Macualay duration in a similar manner.

Our earlier estimate of DV01 can be reconciled with our duration measure as follows. DV01 is the dollar value of one basis point. Let's multiply this by 100 to get the dollar value of point (i.e., 1%), $\$0.0209 \times 100 = \$2.09$. If we divide this by the price of the bond, \$108.38, we get an estimate of the percent change in the price of the bond for a one percent change in the yield, $2.09 \div 108.38 = 0.0192$ or 1.92% - similar to our modified duration estimate. Mathematically,

$$\frac{DV01 \times 100}{Price} \times 100 \approx \text{Modified Duration}$$

7.4.4 Bond Portfolio Duration

A convenient feature of Duration is that computing it for a portfolio of bonds is straightforward. Equation 7.6 details how.

$$\text{Portfolio duration} = \underbrace{\frac{Value_1}{Value_P}}_{weight_1} \times D_1 + \underbrace{\frac{Value_2}{Value_P}}_{weight_2} \times D_2 + ... + \underbrace{\frac{Value_N}{Value_P}}_{weight_N} \times D_N \tag{7.6}$$

$Value_1$ through $Value_N$ correspond to the market values of the N bonds in the portfolio. $Value_P$ is the total market value of the portfolio. The ratios of bond values to portfolio value

[7] Modified 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.

Copyright © 2020 Michael R. Roberts

are referred to as **value-weights**. D_1 through D_N are the N individual bond durations. Equation 7.6 shows that the duration of a bond portfolio is a value-weighted average of component bond durations.

For example, consider the bond portfolio in table 7 consisting of four bonds, A through D. The market value of each bond held in the portfolio is listed in the Value column. The weight for each bond is the corresponding market value divided by the portfolio market value. For example, the weight for bond A is $100 \div 600 = 0.1667$. The duration for the portfolio is the weighted sum of the individual bond durations.

$$\begin{aligned} \text{Portfolio duration} &= 0.1667(0.50) + 0.3333(1.86) + 0.0833(12.50) + 0.4167(9.64) \\ &= 5.76 \end{aligned}$$

Bond	Value ($mil)	Weight (%)	Duration
A	100	16.67	0.50
B	200	33.33	1.86
C	50	8.33	12.50
D	250	41.67	9.64
Portfolio	600	100.00	5.76

Table 7: Portfolio Duration

Even if the composition of the portfolio doesn't change over time - no buying or selling of bonds - the portfolio duration will change as the bonds get closer to their maturity or as their market values fluctuate.

7.5 Application: Hedging a Bond Portfolio

How can we hedge, or eliminate, interest rate risk from our investments? Let's see with an example. Table 8 provides information for three hypothetical bonds as of January 19, 2022. The letters in parentheses correspond to the stock ticker symbols of each company. The maturity column provides the date on which the bonds mature.

7.5.1 Interest Rate Exposure

Let's start by computing the yields and DV01 for all three bonds, even though we only own the first two. To do so, we need to know that the bonds make semi-annual coupon payments

Copyright © 2020 Michael R. Roberts

Bonds	Face Value ($000s)	Coupon Rate (%)	Market Value ($000s)	Maturity
Walt Disney Co. (DIS)	11,032	3.375	12,193	1/19/27
McDonald's Corp. (MCD)	13,255	2.125	13,490	1/19/30
Unilever Capital Corp. (UL)	23,000	1.750	21,858	1/19/32

Table 8: Bond Information

beginning 6 months from today. The yields can be found by solving the fundamental value relations of each bond for y using a computer.

$$\begin{aligned} 12,193,000 &= \frac{186,165}{(1+y/2)} + \frac{186,165}{(1+y/2)^2} + ... + \frac{186,165 + 11,032,000}{(1+y/2)^{10}} \implies y^{DIS} = 1.20\% \\ 13,490,000 &= \frac{140,834}{(1+y/2)} + \frac{140,834}{(1+y/2)^2} + ... + \frac{140,834 + 13,255,000}{(1+y/2)^{16}} \implies y^{MCD} = 1.89\% \\ 21,858,000 &= \frac{201,250}{(1+y/2)} + \frac{201,250}{(1+y/2)^2} + ... + \frac{201,250 + 23,000,000}{(1+y/2)^{20}} \implies y^{UL} = 2.31\% \end{aligned}$$

The DV01 for each bond starts by repricing the bond at its current yield less one basis point.

$$\begin{aligned} \text{DIS: } & \frac{186,165}{(1+0.0119/2)} + \frac{186,165}{(1+0.0119/2)^2} + ... + \frac{186,165 + 11,032,000}{(1+0.0119/2)^{10}} = 12,198,654 \\ \text{MCD: } & \frac{140,834}{(1+0.0188/2)} + \frac{140,834}{(1+0.0188/2)^2} + ... + \frac{140,834 + 13,255,000}{(1+0.0188/2)^{16}} = 13,499,901 \\ \text{UL: } & \frac{201,250}{(1+0.0230/2)} + \frac{201,250}{(1+0.0230/2)^2} + ... + \frac{201,250 + 23,000,000}{(1+0.0230/2)^{20}} = 21,877,872 \end{aligned}$$

The DV01s are

$$\begin{aligned} DV01^{DIS} &= 12,198,654 - 12,193,000 = \$5,654 \text{ (\$0.0513 per \$100 of face value),} \\ DV01^{MCD} &= 13,499,901 - 13,490,000 = \$9,901 \text{ (\$0.0747 per \$100 of face value), and} \\ DV01^{UL} &= 21,877,872 - 21,858,000 = \$19,872 \text{ (\$0.0864 per \$100 of face value).} \end{aligned}$$

(The per \$100 of face value figures are computed by dividing $DV01$ by the face value of the bond and multiplying by 100.) These calculations show that a one basis point increase in interest rates leads to a \$15,555 (5,654 + 9,901) decline in the value of our bond portfolio, or \$0.0640 decline per \$100 of face value.

Copyright © 2020 Michael R. Roberts

7.5.2 Hedge Position

When interest rates increase, the value of all three bonds will fall, though by different amounts, according to our $DV01$ estimates. As just shown, our portfolio value will fall by \$0.640 per \$100 of face value. We need to take a position in the Unilever bond that offsets this decline. In other words, when interest rates increase, our position in the Unilever bond increases.

We can accomplish this is by taking a **short position**, i.e., having short-sold, in the Unilever bond. Remember, short-selling an asset means borrowing it and then selling it. However, eventually we have to buy the asset back and return it to whomever we borrowed it. If the price of the asset falls after having sold it, this is good for our short position because we can buy it back on the cheap.

The question is: How much of the Unilever bond should we short so that changes in the value in our **long position**, i.e., the assets we bought, are exactly offset by changes in our short position? To answer this, it helps to know that the $DV01$ of a portfolio containing N bonds is just a weighed sum of the individual bond $DV01$s.

$$DV01 = \frac{Face_1}{100} \times DV01_1 + \frac{Face_2}{100} \times DV01_2 + ... + \frac{Face_N}{100} \times DV01_N \tag{7.7}$$

where $Face_i$ is the face value of bond i and $DV01_i$ is the $DV01$ for bond i expressed per \$100 for $i = 1, ..., N$. In our example, the $DV01$ for the portfolio containing *all three* bonds is

$$\begin{aligned} DV01 &= \frac{Face_{DIS}}{100} \times DV01_{DIS} + \frac{Face_{MCD}}{100} \times DV01_{MCD} + \frac{Face_{UL}}{100} \times DV01_{UL} \\ &= \frac{11,032,000}{100} \times 0.0513 + \frac{13,255,000}{100} \times 0.0747 - \frac{Face_{UL}}{100} \times 0.0864 \\ &= 15,555 - \frac{Face_{UL}}{100} \times 0.0864, \end{aligned}$$

where the minus sign in front of the UL term indicates our short position.

To perfectly hedge interest rate risk, we want the $DV01$ of our portfolio to be zero. A zero DV01 means that a one basis point decrease in interest rates changes our portfolio's value by zero, i.e., no change. This is a **perfect hedge**. Setting the left side of the expression above to zero and solving for $Face_{UL}$ tells us how much face value of the Unilever bond we should short to ensure a perfectly hedged portfolio.

$$\begin{aligned} 0 &= 15,555.32 - \frac{Face_{UL}}{100} \times 0.0864 \\ \Longrightarrow \quad Face_{UL} &= \frac{15,555.32 \times 100}{0.0864} = \$18,003,809 \end{aligned}$$

Copyright © 2020 Michael R. Roberts

We need to short \$18,003,809 of face value of Unilever bonds to perfectly hedge our bond portfolio against interest rate risk.

If instead we wanted to maintain some interest rate risk exposure, we could replace zero in the equation above with the amount of exposure with which we are comfortable. For example, if we were willing to accept a \$5,000 swing in our portfolio value for a 1 basis point change in interest rates, then we only need short

$$
\begin{aligned}
& 5,000 = 15,555.32 - \frac{Face_{UL}}{100} \times 0.0864 \\
\Longrightarrow \quad & Face_{UL} = \frac{(15,555.32 - 5,000) \times 100}{0.0864} = \$12,216,783
\end{aligned}
$$

of face value of Unilever bonds.

As interest rates and our risk exposure change over time, we'll need to update our hedge by buying and selling bonds to ensure our portfolio's DV01 is zero. In practice, the maintenance of the hedge is infrequent beacuse of transaction costs (e.g., bid-ask spread, brokerage fees). Most hedging practices rely on the portfolio $DV01$ crossing a threshold - positive and negative - before rebalancing the hedge.

7.6 Application: Immunizing a Bond Portfolio

Immunization locks in a value for a bond portfolio at a point in the future. That point could be defined by an investment or planning period, e.g., 1-, 2-, 5-years. Regardless, by structuring the portfolio so that its duration is equal to the time horizon, we can ensure a specific value for our portfolio at that future point in time...pretty cool. Let's see how this works with an example.

Consider our portfolio of Disney and McDonald's bonds. We saw above how shorting the right amount of Unilever bonds we could hedge the interest rate risk of our portfolio. Now let's determine how we can rebalance our portfolio - change the relative amounts of Disney and McDonald's bonds - to immunize the portfolio at a 5-year horizon. That is, let's ensure the value of our portfolio five years from today.

The market values of Disney and McDonald's bonds are, from Table 8, \$12.193 and \$13.490 million, respectively. Our portfolio's market value is the sum of these values, \$25.683 million. The bond durations are 4.66 and 7.41, respectively. (See Table 6 for a reminder of how these estimates are computed.) According to equation 7.6, the duration of our portfolio

Copyright © 2020 Michael R. Roberts

is

$$\text{Portfolio duration} = \frac{12,193,000}{25,683,000} \times 4.66 + \frac{13,490,000}{25,683,000} \times 7.41 = 6.10$$

We want our portfolio duration to be 5, so we need to rebalance our portfolio or take a position in another bond. Let's focus on the former. What we want to find are the **portfolio weights** dictating the fraction of portfolio value allocated to each bond such that the portfolio duration is equal to our target, 5. Mathematically,

$$\text{Target duration} = 5 = w_{DIS} \times 4.66 + (1 - w_{DIS}) \times 7.41$$

where w_{DIS} is the value-weight of the Disney bond, and $(1 - w_{DIS}) = w_{MCD}$ is the value-weight of the McDonald's bond. (The weights sum to one.). Solving for w_{DIS} produces

$$\begin{aligned} & w_{DIS} = \frac{5 - 7.41}{4.66 - 7.41} = 0.8764 \\ \implies \quad & w_{MCD} = (1 - w_{DIS}) = 0.1236. \end{aligned}$$

To hit our target portfolio duration of 5 years, we need 87.64% of our portfolio invested in Disney bonds and 12.36% of our portfolio invested in McDonald's bonds.

In dollar terms, these weights imply our target portfolio contains $22,534,597 in Disney bonds and $3,148,403 in McDonald's bonds. To get to this target, we must purchase $10,341,597 worth of Disney bonds by selling the same dollar amount of McDonald's bonds. In face values, this means buying $9,356,885 of Disney face value and selling $10,161,443 of McDonald's face value.

We need to rebalance our portfolio to maintain the immunization. In other words, as time evolves and we approach our target end date, the duration of our portfolio must change to match the remaining time. Part of this change occurs naturally through the reduction in the maturity of the bonds. The rest comes from buying and selling appropriate amounts of the bond, similar to what we just computed. However, because of transaction costs, rebalancing occurs infrequently, much like hedging.

7.7 Default Risk

When there is a chance the borrower can't repay the loan, the future cash flows become risky. We address this risk by using *expected* cash flows, just like we used expected prices to compute expected returns.

Copyright © 2020 Michael R. Roberts

Consider a simple example. We have a one year zero coupon bond currently priced at \$92 per \$100 of par value. Assume there is a 10% probability that the borrower will default and not 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× 100 = \$60. (Bondholders typically have some recourse against the borrower's assets when they default.)

The yield to maturity of this bond is computed using the *promised* cash flow, \$100.

$$Price_0 = \frac{CF_1}{(1+y)} \implies y = \frac{CF_1}{Price_0} - 1 = \frac{100}{92} - 1 = 0.087.$$

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 then sum. In our bond example, the expected cash flow one year from today is

$$\begin{aligned} E(CF_1) &= Pr(\text{Default}) \times \text{Cash flow in default} + Pr(\text{No default}) \times \text{Cash flow no default} \\ &= 0.10 \times 60 + (1 - 0.10) \times 100 \\ &= \$96. \end{aligned}$$

More generally, an expected value is just a weighted sum in which the weights are probabilities. So for a random variable, X, that can take on values $x_1, x_2, ..., x_N$ with probabilities $pr_1, pr_2, ..., pr_N$, the expected value of X is:

$$\mathbb{E}(x) = pr_1 x_1 + pr_2 x_2 + ... + pr_N x_N \tag{7.8}$$

When the variable can take on a continuum of values, like the value of a company or the price of a stock, the intuition remains the same but we need a little calculus. See the technical appendix for details.

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.

$$r = \frac{E(CF_1)}{Price_0} - 1 = \frac{96}{92} - 1 = 0.043$$

Copyright © 2020 Michael R. Roberts

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. That overstates the cost of capital because it ignores default risk. The debt cost of capital is the expected return on the company's debt and, unless the company is financially sound, the yield to maturity and expected return can be quite different as our numerical example showed.

7.7.1 Credit Ratings

We saw above that computing interest rate risk is relatively easy. DV01 and duration provide quick and accurate estimates of the change in value of our bond or bond portfolio that occurs when interest rates change. Estimating default risk is quite a bit more difficult. We have to understand the financial risk of the issuing entity (e.g., company or government), as well as its incentives, to understand the likelihood that it may default at some point during the life of the bond.

Fortunately, there are rating agencies whose primary job is to estimate those likelihoods. Companies like Moody's, Standard and Poors (S&P), and Fitch distill bond default information into **credit ratings** that characterize the default risk investors face when purchasing bonds. While there are other ratings agencies, these three are the largest and cover the largest number of bonds.

Rating scales are depicted in table ??, which is adapted from Wikipedia.[8] 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.

[8] The ratings from Fitch overlap with those from S&P except that there is no "SD" in the Fitch rating scale.

Copyright © 2020 Michael R. Roberts

Moody's	S&P and Fitch	Average Default Probability	Description
		Investment Grade	
Aaa	AAA		Prime
Aa1	AA+		
Aa2	AA	0.1%	High grade
Aa3	AA-		
A1	A+		
A2	A	0.25%	Upper medium grade
A3	A-		
Baa1	BBB+		
Baa2	BBB	1.0%	Lower medium grade
Baa3	BBB-		
		Below Investment Grade	
Ba1	BB+		
Ba2	BB	7.5%	Non-investment grade
Ba3	BB-		
B1	B+		
B2	B	20%	Highly speculative
B3	B-		
Caa1	CCC+		
Caa2	CCC	34%	Substantial risk
Caa3	CCC-		
Ca	CC		Extremely speculative
	C		Default is imminent
C	RD		
	SD		In default
	D		

Table 9: Credit Ratings Scales

While useful, credit ratings have their limitations. They can be slow to incorporate new information. They can also prove quite inaccurate at times, as they did for many securitizations that defaulted during the 2008 financial crisis. 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.

Copyright © 2020 Michael R. Roberts

7.8 Inflation Risk

In addition to interest rate and credit risk, most bonds 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 many goods and services because the price is going up more quickly than we are earning interest.

7.8.1 Application: TIPS

In 1997, **Treasury Inflation Protected Securities** or (**TIPS**) were introduced to address the shortcomings of **nominal bonds**, or fixed income securities that are not insulated from inflation risk. In a TIPS, the coupon rate is fixed and coupons are paid semi-annually, just like Treasury notes and bonds. However, the principal of TIPS is tied to the prices of goods and services in the economy.[9] 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.

Table 10 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. The CPI used for adjusting the principal as of the dated date was 273.2577 and is referred to as the **reference CPI**.[10] 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 $1,000 \times 0.000125/2 = \$0.625$ 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.)

[9] More precisely, the principal of the securities is indexed to the **Consumer Price Index for All Urban Consumers** (**CPI-U**).

[10] A 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.

Copyright © 2020 Michael R. Roberts

					Hypothetical	
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 10: 5-year, 0.125% Coupon TIPS Example Cash Flows

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, 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 $282.3464 \div 273.2577 - 1 = 0.03326$, 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* coupon 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.

Copyright © 2020 Michael R. Roberts

7.9 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 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. This is the same thing as the internal rate of return.
- The yield curve is a visual representation of the term structure of interest rates. For each term structure, there is a yield curve.
- Realized returns are the returns we as investors actually earn. Expected returns are the returns we expect (or hope) to earn. 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. When relative mispricings in financial markets arise, investors are quick to pounce on arbitrage opportunities. But, by doing so, the relative prices come back in line.

  Take our T-note example. As more people buy T-notes, the price will increase. Likewise, as more people sell zero coupon bonds, their prices will fall. As a result of these two forces, the relative prices of the T-notes and zero coupon bonds will get closer together, thereby eliminating the arbitrage opportunity. Another way to view this phenomenon is that most financial markets are extremeley competitive, and this competition makes sure prices do not get too far away from their fundamentals and if they do it's not for long.

  This is *not* to say that market prices are never wrong. 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**

Copyright © 2020 Michael R. Roberts

of arbitrage.[11] 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 dollar duration, Macaulay duration, and modified duration. Default risk is comprised of two components: the probability of default and the loss given default.

7.10 Technical Appendix

7.10.1 Expected Returns for Defaultable Bonds

To compute the expected return to a defaultable bond, we need to solve for r in our fundamental value relation in which we replace the *promised* cash flows of the bond with the *expected* cash flows.

$$Price = \frac{\mathbb{E}(CF_1)}{1+r} + \frac{\mathbb{E}(CF_2)}{(1+r)^2} + ... + \frac{\mathbb{E}(CF_T)}{(1+r)^T}$$

Solving for r is easy in a spreadsheet program. The challenging part is estimating the expected cash flows, $\mathbb{E}(CF_t)$. Let's illustrate the process with an example.

Consider a bond with a 5% annual coupon, par value of \$1,000, maturing five years from today, and priced at \$854. The yield to maturity, y, can be calculated using the *promised* cash flows.

$$Price = \frac{50}{1+y} + \frac{50}{(1+y)^2} + \frac{50}{(1+y)^3} + \frac{50}{(1+y)^4} + \frac{1,050}{(1+y)^5} \implies y = 0.0873$$

Because the borrower may default, the promised and expected cash flows are not the same. Consequently, the yield to maturity and expected return are not the same. To estimate the expected cash flows, we need to know the probability of default at each point in time over the life of the bond and the recovery rate in default.

Figure 7.10 visualizes the valuation problem in the form of a **decision tree**. The bond price today is \$854, depicted as a cash outflow in period 0. In one year's time, one of two things can happen. The borrower can default, in which case investors receive 60% of (1,000+50), or \$630. We've assumed a recovery rate of 60%. Or, the borrower can make

[11]This term was coined by Andrei Shleifer and Robert Vishny in their 1997 Journal of Finance article by the same name.

Copyright © 2020 Michael R. Roberts

the first coupon payment of \$50.[12] The probability of default is 5.41%, the probability of no default is 94.59%, as indicated by the red numbers in the tree. These probabilities come from Moody's 2006 research report - more recent reports are proprietary - and correspond to the default probabilities of a B-rated bond. The expected cash flow one year from today is

$$\mathbb{E}(CF_1) = \underbrace{0.9549}_{Pr(\text{No default})} \times \overbrace{(50)}^{\substack{\text{Payment}\\\text{when}\\\text{no default}}} + \underbrace{0.0541}_{Pr(\text{Default})} \times \overbrace{(630)}^{\substack{\text{Payment}\\\text{when}\\\text{default}}} = \$81.38.$$

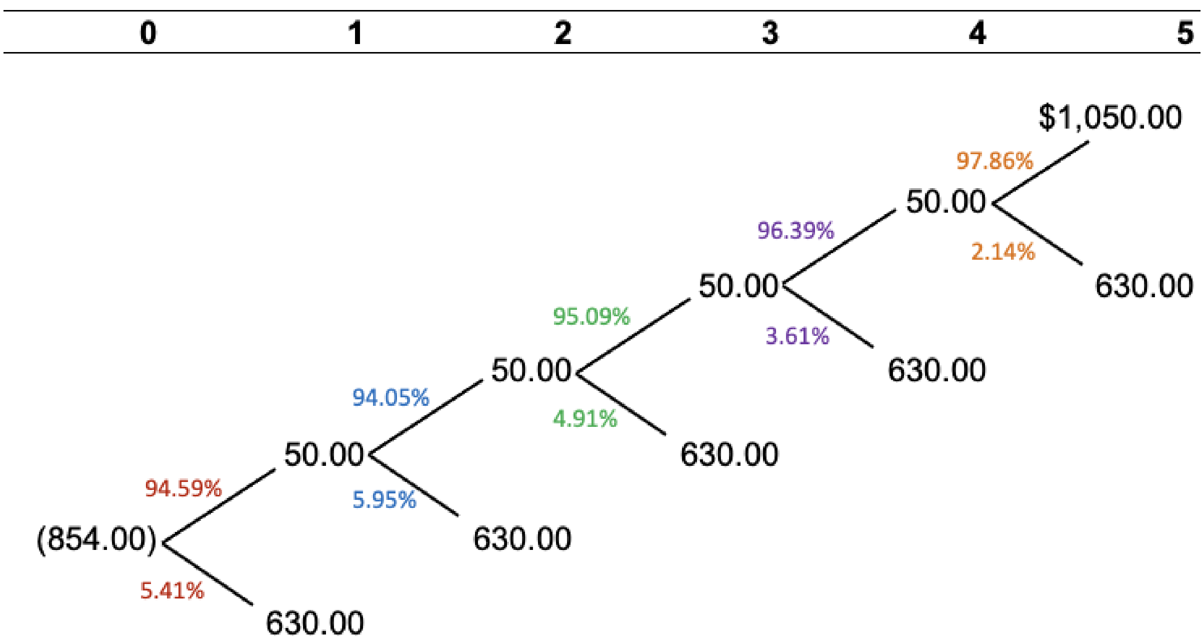

Figure 7.10: Expected Return on a B-rated Bond

If the firm does not default one year from today, then two years from today we will receive another \$50 if the firm doesn't default or we will receive \$630 if it does. The expected cash flow two years from today is

$$\mathbb{E}(CF_2) = \underbrace{0.9459}_{Pr(\text{No default t=1})} \times \left[\underbrace{0.9405}_{\substack{Pr(\text{No default t=2 }|\\\text{No default t=1})}} \times \overbrace{(50)}^{\substack{\text{Payment}\\\text{when}\\\text{no default}}} + \underbrace{0.0595}_{\substack{Pr(\text{Default t=2 }|\\\text{No default t=1})}} \times \overbrace{(630)}^{\substack{\text{Payment}\\\text{when}\\\text{default}}} \right] = \$79.95$$

[12]it might seem odd that a borrower could default and pay more than what they would pay if they didn't default. But, don't view this recovery amount as cash lying around the firm that managers could pay. Instead, think of it as what the creditors would get from selling off the assets (Chapter 7) or the value of their new claims in reorganization (Chapter 11).

Copyright © 2020 Michael R. Roberts

Notice we multipy the weighted average of the cash flows by the probability of not defaulting in the first period (0.9459), which must happen in order for the investor to arrive at period 2.

Continuing this logic gives us the expected cash flows for period three through five.

$$\begin{aligned}
\mathbb{E}(CF_3) &= 0.9459 \times 0.9405 \times [0.9509(50) + 0.0491(630)] = 69.83 \\
\mathbb{E}(CF_4) &= 0.9459 \times 0.9405 \times 0.9509 \times [0.9639(50) + 0.0361(630)] = 59.99 \\
\mathbb{E}(CF_5) &= 0.9459 \times 0.9405 \times 0.9509 \times 0.9639\,[0.9786(50) + 0.0214(630)] = 846.93
\end{aligned}$$

With the expected cash flows, we can compute the expected return on the bond by solving for r in our fundamental value relation.

$$Price = \frac{81.38}{1+r} + \frac{79.95}{(1+r)^2} + \frac{69.83}{(1+r)^2} + \frac{59.99}{(1+r)^2} + \frac{846.93}{(1+r)^2} \implies r = 0.0696$$

The expected return, 6.96%, is less than the yield-to-maturity, 8.73%. Notice that the first four expected cash flows are all *greater* than the promised cash flows because the former include the possibility that the borrower defaults, and we receive the recovery value of the bond, \$630. However, the last expected cash flow containing the principal repayment is significantly less than the promised cash flow, which is what ultimately causes the expected return to be less than the yield.

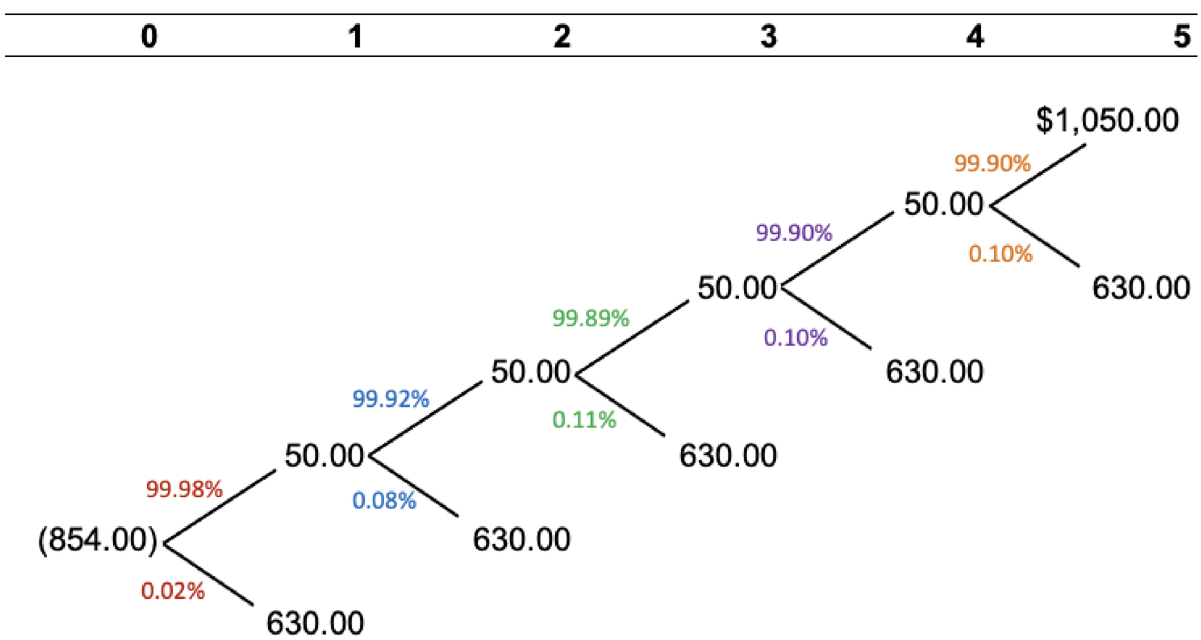

Figure 7.11: Expected Return on an A-Rated Bond

Copyright © 2020 Michael R. Roberts

Figure 7.11 presents a similar analysis for an A-rated bond. Notice that the default probabilities are significantly lower than those for the b-rate bond. Consequently, the expected cash flows are closer to the promised cash flows, and the expected return closer to the yield to maturity.

$$Price = \frac{50.12}{1+r} + \frac{50.45}{(1+r)^2} + \frac{50.59}{(1+r)^2} + \frac{50.48}{(1+r)^2} + \frac{1046.33}{(1+r)^2} \implies r = 0.0869$$

The 8.69% expected return is only slightly less than the yield to maturity because of the negligible likelihood of default. Even if there was no recovery in default - a total loss - the expected return on the A-rated bond would be 8.64%.

7.10.2 Taxes on Bond Returns

Recall the bond considered at the beginning of this chapter - a 10-year, zero-coupon bond with a face value of $1 million and 4% semi-annual compounded APR. The value today was shown to be $672,971.33. That the price is below the face value at issuance is referred to an **original-issue discount** or **OID**.

The lack of interest payments could create a tax preference for zero coupon bonds if they were 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 OID is a way for the bond to implicitly earn interest over its term. Thus, the tax authorities tax OID bonds in the following manner.

Take our 10-year 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

$$P_1 = \frac{1,000,000}{(1+0.03/2)^{18}} = \$764,911.59.$$

Copyright © 2020 Michael R. Roberts

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.

7.10.3 Proofs

Inverse Relation Between Bond Prices and Interest Rates

This section contains more formal proofs of several results beginning with the inverse relation between bond prices and interest rates. Taking the derivative of our fundamental value relation produces the following expression.

$$\frac{\partial P_t}{\partial r} = -\frac{CF_{t+1}}{(1+r)^2} - 2\frac{CF_{t+2}}{(1+r)^3} - ... - T\frac{CF_T}{(1+r)^{T+1}}$$

This expression is less than zero, implying an inverse relation. More generally, this result shows that prices and expected returns are inversely related. As expected returns, r, increase, prices decline.

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 is a Semi-Elasticity

Start with the basic pricing formula using the annual yield on the bond.

$$P_t = \frac{CF_{t+1}}{(1+y/k)} + \frac{CF_{t+2}}{(1+y/k)^2} + ... + + \frac{CF_{t+T}}{(1+y/k)^T}$$

Copyright © 2020 Michael R. Roberts

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.

$$\begin{aligned}\frac{1}{P_t}\frac{\partial P_t}{\partial y} &= -\frac{1}{1+y/k}\left[\frac{1}{k}\left(\frac{PV(CF_{t+1})}{P_t}\right) + \frac{2}{k}\left(\frac{PV(CF_{t+2})}{P_t}\right) + \dots + \frac{T}{k}\left(\frac{PV(CF_{t+T})}{P_t}\right)\right] \\ &= -\frac{1}{1+y/k}\sum_{s=1}^{T} \underbrace{\left(\frac{s}{k}\right)}_{\text{Year of } s^{th} \text{ cash flow}} \times \underbrace{\left(\frac{PV(CF_{t+s})}{P_t}\right)}_{\text{Weight of } s^{th} \text{ cash flow}}\end{aligned}$$

where $PV(CF_s)$ is the present value of the cash flow arriving in period s.

7.11 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 (*Coupon bond yields*) 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?

7.4 (*Coupon bond valuation*) You are offered an 8% coupond bond with a face value of \$1,000,000. The bond has three years to maturity and the coupons are paid semi-annually. If the yield-to-maturity is 10%, what is the fair price of the bond?

7.5 (*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	?

Copyright © 2020 Michael R. Roberts

The coupon rate and yield are APRs. Coupons and interest are paid and compounded annually.

Using this information, answer the following questions.

(a) What are the values for each "?" in the table?

(b) Is there an arbitrage opportunity? If so, how would you take advantage of it? Be clear to identify the positions and corresponding cash flows.

7.6 (*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.

a. Assuming the federal government does not default on the bond, what will be the realized return on the bond when it matures?

b. What is the annual yield on of the bond?

c. How are the realized return and annual yield for this bond related?

d. If immediately after purchasing your bond the annual 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?

7.7 (*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:

(a) one-year zero-coupon bonds, and

(b) 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.

a. How much money must you put aside today assuming that interest rates stay unchanged?

Copyright © 2020 Michael R. Roberts

b. Suppose you invest the money you need to set aside today - your answer to the previous 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.

c. 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, wow much extra money will have left over, or will you have enough to make the second payment?

7.8 (*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?

7.9 (*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.

Copyright © 2020 Michael R. Roberts

7.10 (*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?

7.11 (*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.

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.)

7.12 (*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?

Copyright © 2020 Michael R. Roberts

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.

7.13 (*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.

7.14 (*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 per $100 of face value if the current annual yield is 10%.

7.15 (*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?

7.16 (*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 (if relevant), and annual compounding.

Copyright © 2020 Michael R. Roberts

Bond	Coupon rate (%)	Maturity (Years)	Price ($)
A	0	1	90
B	0.10	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?

7.17 (*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?

7.18 (*Real and nominal bond returns*) TODO

7.19 (*Default risk*) Torvax Inc. recently issued a Ba-rated, seven-year, annual coupon bond with an 8% coupon rate. According to Moody's, the average cumulative default probabilities for such a bond are as follows.

Year	1	2	3	4	5	6	7
Probability	1.18	2.98	4.85	6.52	7.86	8.9	9.68

The expected recovery rate in default is 52%, applied to both interest and principal.

Using this information, answer the following questions.

a. What is the price of the bond for $1,000 of face value?

b. How money ($) is recovered by investors if the bond defaults?

c. What are the expected cash flows for each period?

d. What is the expected return on the bond?

7.20 (*Default risk probability concepts*) The cumulative default probabilities are equal to one minus the **survivor function**, $S(t)$, which is the probability of the firm defaulting after time t, $Pr(Default \geq t)$. The **hazard function**, $h(t)$, is the probability that the firm defaults at time t conditional on it not having defaulted before that point, $Pr(Default = t | Default \geq t)$. Finally, the **marginal probability**, $f(t)$, the unconditional probability that the firm defaults at time t, $Pr(Default = t)$. Mathematically, these functions are related as follows.

Copyright © 2020 Michael R. Roberts

$$h(t) = \frac{S(t) - S(t-1)}{S(t-1)} = \frac{f(t)}{S(t-1)}$$

Answer the following questions.

a. Using the the cumulative default probabilities from the previous problem, what is the hazard function? That is, estimate $h(t)$ for each $t = 1, ..., 7$. Plot your results. What does the shape of the hazard function tell you about the likelihood of default over the term of the bond?

b. Using the the cumulative default probabilities from the previous problem, what are the marginal probabilities for $t = 1, ..., 7$?

c. Compute the internal rate of return on the bond for each default scenario - years one through seven - and for the no default scenario. Take a probability weighted average of these yields where the probabilities are the marginal probabilities, $f(t)$, you computed in the previous question. What value do you get? What does this number represent? How does it compare to the expected return?

7.21 (*Measuring default risk*) Altman's Z-score is popular measure characterizing the default risk of a company. It is defined as follows.

$$\begin{aligned} Z &= 3.10\frac{EBIT}{TotalAssets} + 1.00\frac{Sales}{TotalAssets} + 0.42\frac{BookEquity}{TotalLiabilities} \\ &+ 0.85\frac{RetainedEarnings}{TotalAssets} + 0.72\frac{WorkingCapital}{TotalAssets} \end{aligned}$$

Does a higher Z-score correspond to a firm with a higher or lower likelihood of financial distress? Explain your answer.

7.22 (*After-tax bond yield comparisons*) Consider the Treasury and Municipal bond yield curves in figure 7.8. 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.

Copyright © 2020 Michael R. Roberts

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?

7.23 (*After-tax bond yield comparisons*) Consider the Treasury and Municipal bond yield curves in figure 7.8.

- 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?

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Chapter 8

Investing: Stocks

Fundamental value relation

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

A portfolio is a collection of assets - stocks, bonds, real estate, etc. What happens to our returns - realized and expected - and risk when we buy more than one stock or bond?

This chapter

- introduces stocks as an investment for savers and source of funding for corporations,
- discusses how to buy and sell stock,
- distinguishes between expected returns, the r in our fundamental value relation and what we expect to earn each period as investors, and realized returns, what we actually earn each period,
- considers the implications of taxes on investor returns,
- shows that stock valuation is a straightforward application of the fundamental value relation, but there are several different approaches to estimating the cash flows to shareholders,
- presents several methods for estimating expected returns and risk of stock investing, and

Copyright © 2020 Michael R. Roberts

- applies our fundamental value relation to answer several questions including:
  - How can we estimate an expected stock return and the risk of investing in a stock?
  - 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 stock?
  - How can we identify whether a stock is over- or under-valued by the market?

8.1 What is Stock?

Stock, or **equity**, is a claim to the cash flows of a firm, just like a loan or bond. In many instances, stock ownership also comes with the right to vote on corporate governance issues. Stockholders are the owners of the company, and the degree of ownership is proportional to the number of shares one holds.

For example, imagine company with 100 shares oustanding. Joe owns 20, Mary owns 80. In this simple example, Joe owns 20% of the company, Mary 80%. Because Mary owns the majority of the stock, she gets to decide what happens with the company. If there were three shareholders each owing one third of the company, then two would have to agree to any changes, thereby ensuring a majority.

Whereas the cash flows received by lenders are in the form of interest and principal, the cash flows received by stockholders are in the form of dividends - cash and sometimes additional shares - and stock buybacks or repurchases. Dividends tend to be more consistent payouts, much like a coupon payment on a bond, whereas repurchases afford companies more flexibility in their timing and magnitude.

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 and many companies, particularly younger and smaller companies, do neither for some time. Though, as we'll see, eventually companies must get money to shareholders if they wish to raise money from them.

Another difference between lenders and stockholders is that the latter only get paid *after* all of the lenders have been paid. That is, stockholders are lower in the **priority structure** of companies' liabilities and, as such, stockholders are referred to as **residual claimants**. They get what's left over.

Stock comes in several different flavors.

Copyright © 2020 Michael R. Roberts

8.1.1 Common Stock

When people talk about "stock," they're often referring to **common stock**. Owning stock is owning a share or portion of a company; hence the name **shareholders**. Common stockholders are at the very bottom of the liability priority structure. They only receive money after every other investors has received what they are owed. Stockholders receive money from the company through two channels: **dividends** from the company's earnings and **stock repurchases** or **buybacks** in which the company buys back stock from investors on the stock market.

Common stock typically confers voting rights - one vote per share - though these rights can vary in **dual-** and **tri-class** share structures. For example, firms in which founders and families play a large role often have dual-class share structures. For example, GoPro (ticker symbol GPRO) has 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 the Nicholas Woodman, the founder, have 10 votes per share. Alphabet (ticker symbol GOOGL) has a tri-class structure in which A class shares issued to the public have one vote per share, B class shares are held by the founders, Sergey Brin and Larry Page, and several directors of the company, and C class shares issued to the public with no voting rights. Ultimately, the variation in voting rights comes down to control. Founders and their families want to retain control of their companies while also reaping the benefits of public equity financing.

In international settings, dual class shares often separate voting rights between domestic and foreign investors. Class A shares issued to domestic investors have voting rights, class B shares issued to foreigners do not.

8.1.2 Preferred Stock

Preferred stock sits above common stock, and below all debt, in the priority structure of corporate liabilities and, in some ways, is more like a bond than stock. Preferred stock comes with a par value and a fixed dividend that operates like a coupon payment. In some cases, preferred stock even comes with a maturity data. If a dividend payment is missed, then dividends accumulate and must be paid out before paying any money to common stockholders. Unlike debt, firms cannot be forced into bankruptcy by preferred shareholders because of a missed dividend payment.

Preferred stock typically does not come with voting rights, though missing a dividend payment can sometimes activate voting rights for these investors. Relative to common stock,

Copyright © 2020 Michael R. Roberts

preferred stock is rare and is typically issued by utilities for regulatory purposes and banks because of the tax treatment of dividends.

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 a T-bill yield.

8.2 Buying and Selling Stock

All companies have stock, but relatively few have stock that is publicly traded. Of the 6.1 million companies in the U.S. in 2019, fewer than 13,000 are listed on a U.S. exchange. When firms do issue stock to the public, they typically work with investment banks who underwrite and help sell the stock to the public in a two step process. First, stock is sold in the **primary market** to institutional investors. Buying and selling (a.k.a., trading) stock after this initial placement occurs in the secondary market.

8.2.1 Price Quotes

Depending on whether we want to buy or sell a stock, we'll do so at different prices. Table 1 shows price quotes for Amazon.com Inc. (AMZN) and Iveric Bio Inc. (ISEE) as of October 21, 2022. The **ask price** is the current price at which we can buy stock. The **bid price** is the current price at which we can sell stock. The difference between the bid and the ask, know as the **bid-ask spread**, is a transaction cost we face when we trade stock (and many other securities).

The size figues are 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 per share. Of course, 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, only 500 shares would be sold at \$117.25. The remaining will likely to sell at a lower price, or we can wait to sell those extra shares at a more favorable price.

Copyright © 2020 Michael R. Roberts

	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**.

Market Order

A **market buy** or **market sell** order requests that the order be executed as quickly as possible, regardless of the execution price. 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 than the $117.26 per share if prices change before the order is received. Or, 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 will be executed is uncertain. For small orders and in normal markets, the price risk is relatively low. But, for particularly volatile stocks or in turbulent markets, prices can quickly change by large amounts.

Limit Order

Limit orders are the complement to market orders. 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 $10 per share when Amazon is currently trading at $117.25, there is

Copyright © 2020 Michael R. Roberts

a good chance that our order will never get executed. Limit orders are good in turbulent markets or when we are particularly sensitive to changes in the price.

In addition to specifying a price for our limit order, we need to specify a term, or for how long the order should be valid. Many limit orders are only valid until the end of the day on which they are submitted. At the end of the day, unfilled limit orders are simply cancelled and may no longer be executed. Other limit orders can be valid for a number of days (7, 30, 60, etc.). Some limit orders are valid until filled, meaning the order must be cancelled at some point or it will remain valid indefinitely.

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. Consider a stop buy order. We might execute a stop buy 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 only 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. Wemight execute 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 order stop sell will convert to a market sell and will be immediately executed.

8.3 Realized Returns

How do we make (or lose) money when we invest in - that is, buy - a stock? Table 2 presents monthly return data from January 2021 to December 2021 for Microsoft. 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 or certain holidays.) The dividend per share column presents the dividend amount that each share receives. Microsoft pays quarterly dividends.

Copyright © 2020 Michael R. Roberts

Date (YYYYMMDD)	Price per Share	Dividend per Share	Price Return	Dividend Yield	Total Return	Annual Return
20201231	222.42					
20210129	231.96		4.29%	0.00%	4.29%	
20210226	232.38	0.56	0.18%	0.24%	0.42%	
20210331	235.77		1.46%	0.00%	1.46%	
20210430	252.18		6.96%	0.00%	6.96%	
20210528	249.68	0.56	(0.99%)	0.22%	(0.77%)	
20210630	270.90		8.50%	0.00%	8.50%	
20210730	284.91		5.17%	0.00%	5.17%	
20210831	301.88	0.56	5.96%	0.20%	6.15%	
20210930	281.92		(6.61%)	0.00%	(6.61%)	
20211029	331.62		17.63%	0.00%	17.63%	
20211130	330.59	0.62	(0.31%)	0.19%	(0.12%)	
20211231	336.32		1.73%	0.00%	1.73%	52.5%

Table 2: Microsoft Monthly Stock Price Data

The monthly return column is computed using 8.2 but without the expectations. In other words, the realized return between $t-1$ and t is

$$\text{Realized return}_{t-1,t} = \underbrace{\frac{P_t - P_{t-1}}{P_{t-1}}}_{\text{Price return}} + \underbrace{\frac{D_t}{P_{t-1}}}_{\text{Dividend yield}} .$$

The table also shows the decomposition of the total return into its price return and dividend yield components. For example, the monthly return for November 2021 - i.e., the end of October to the end of November - is

$$\frac{330.59 - 331.62}{331.62} + \frac{0.62}{331.62} = -0.0031 + 0.0019 = -0.0012.$$

The final column shows the annual compounded return for 2021 by compounding the monthly returns.

$$\begin{aligned}\text{Annual return}_{t-12,t} &= (1 + r_{t-12,t-11}) \times (1 + r_{t-11,t-10}) \times ... \times (1 + r_{t-1,t}) - 1 \\ &= (1 + 0.0429) \times (1 + 0.0042) \times ... \times (1 + 0.0173) \\ &= 0.525\end{aligned}$$

For each dollar we had invested in Microsoft stock at the start of 2021, we earned \$0.525.

Copyright © 2020 Michael R. Roberts

The equation we used to get the annual return might look strange. In fact, it's not new. It's just our effective annual rate (EAR) formula (equation 3.3) modified to allow the periodic rate to vary over time. Remember, the EAR from chapter (3) is

$$r = (1 + i)^k - 1,$$

where i is the periodic interest rate and k is the number of periods in a year (12 in this case because of monthly returns). But, our EAR expression can also be written just like our annual return expression,

$$r = (1 + i) \times (1 + i) \times ... \times (1 + i) - 1.$$

When we introduced the EAR, the periodic interest rate, i, was the same number every period. This allowed us to write it 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.

8.3.1 Dividends

Let's clarify how dividends work in practice. A dividend begins with a declaration by the company that it will pay a dividend in the future. The date on which this declaration occurs is called the **declaration date**. When declaring the dividend, the company set a **record date**, which is a future date on which an investor must be registered as a shareholder to receive the upcoming dividend. The record date in turn determines the **ex-dividend date** or **ex-date**, which usually occurs one business day after the record date. An investor must have purchased the stock *before* to the ex-date in order to receive the upcoming 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 8.1 illustrates the dividend payment process a dividend paid by Microsoft in 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, or 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.[1] An exception occurs for especially large dividends, 25% or more of the share price. In this case, 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.

[1] Notice 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.

Copyright © 2020 Michael R. Roberts

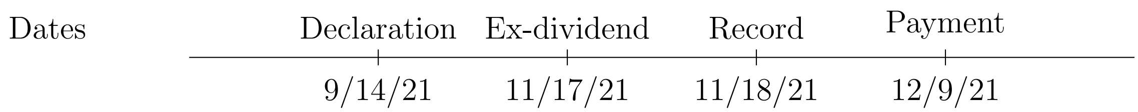

Figure 8.1: Dividend Dates - Microsoft (MSFT)

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).

8.3.2 Stock Splits

Occasionally, a company will engage in a **stock split** to increase the number of shares outstanding without raising any money from the new shares. The reasons for doing is often attributed to a desire to reduce the price of a stock that has gotten so high that retail investors may not be able to purchase a single share.[2] For example, Berkshire Hathaway, the company owned by famed investor Warren Buffest, has class A shares that trade for over $460,000 per share as of December 2022. Few people can afford to buy a 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, meaning each share was converted into three shares. This increase in shares can be seen in the table. 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, $970.50 \div 3 = \$323.50$. 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.

What's important from an investor standpoint is recognizing that we as 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

[2]This rationale is less relevant today when investors can purchase **fractional shares** and therefore are no longer constrained to purchase individual shares.

Copyright © 2020 Michael R. Roberts

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)

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.

$$\text{Return Adj.} = \frac{330.07 \times 3}{970.50} - 1 = 0.0203$$

Companies can 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: Stock Split - Intuitive Surgical (ISRG)

Splits and reverse splits are infrequent but common events for many companies' stocks. The key lessons from our discussion are that they have no effect on investors wealth or returns, once returns are adjusted to reflect the split. So, that 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.

Copyright © 2020 Michael R. Roberts

8.3.3 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 5 presents the capital gains tax rates as of 2021, which are significantly higher than the corresponding income tax rates.

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 5: 2021 Capital Gains Tax Rates by Filing Status and Income Bracket

There is one imporatnt 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.

8.4 Short-term vs. Long-term Investors

Does it matter for value whether one is a short-term or long-term investor? Consider selling a stock one period (day, month, year, etc.) after purchase. According to our fundamental value relation, the purchase price of the stock should equal the sum of the discounted future cash flows. Mathematically,

$$P_0 = \frac{\mathbb{E}(D_1 + P_1)}{(1 + r^E)}, \tag{8.1}$$

where P_0 is the purchase price today, D_1 is the dividend received at the end of the period (if any), and P_1 is the sale price. Because we don't know what the dividend or sale price will be, we use their statistical expectation as our best guess, which we then discount by $(1 + r^E)$.

Copyright © 2020 Michael R. Roberts

Some comments:

- 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**. Had we sold it just before the dividend, we would say **cum-dividend**.

  This assumption is only meant to make our lives a little easier by avoiding having to discount the dividend and the sale price by a different number of time periods. If the stock doesn't pay a dividend before we sell it, then $D_1 = 0$.

- r^E in equation 8.1 is the expected total return, or simply the expected return, of our stock. (The "E" superscript is for Equity.) This variable also goes by the **equity cost of capital** and **levered cost of capital**. A little algebra reveals

$$r^E = \underbrace{\frac{\mathbb{E}(P_1 - P_0)}{P_0}}_{\text{Price return}} + \underbrace{\frac{\mathbb{E}(D_1)}{P_0}}_{\text{Dividend yield}} . \tag{8.2}$$

  The expected return investors earn on a stock is the sum of two components: (i) the price return (a.k.a., **price appreciation/depreciation**, **capital gains/losses**), and (ii) the **dividend yield**.

- The inverse of the dividend yield, $P_0/\mathbb{E}(D_1)$, is the **price-dividend ratio**, which measures how much the stock market is willing to pay for a dollar of dividends. To be precise, the dividend yield in equation 8.2 is the **prospective** or **forward** dividend yield. The ratio of the current dividend, D_0, to the current price is referred to as the **trailing** or **historical** dividend yield.

Equation 8.1 looks a lot different than our fundamental valuation relation. In fact, once we recognize that dividends are how companies distribute money to their shareholders, we'll see that they're identical! That is, from the shareholder's perspective, the cash flows to their claim on the company are the dividends.

Let's use the same logic that we used to write equation 8.1 to write the price one period from today, P_1.

$$P_1 = \frac{\mathbb{E}(D_2 + P_2)}{(1 + r^E)}$$

This equation just says that to whomever we sell the stock one period from today, the price they should pay us, P_1, is equal to the present value of the future cash flows they receive.

Copyright © 2020 Michael R. Roberts

Let's plug this expression for P_1 into equation 8.1.

$$P_0 = \frac{\mathbb{E}\left(D_1 + \frac{\mathbb{E}(D_2+P_2)}{(1+r^E)}\right)}{(1+r^E)} = \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}(D_2+P_2)}{(1+r^E)^2}$$

If we repeat this process for P_2, we get

$$P_0 = \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}\left(D_2 + \frac{\mathbb{E}(D_3+P_3)}{(1+r^E)}\right)}{(1+r^E)^2} = \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}(D_2)}{(1+r^E)^2} + \frac{\mathbb{E}(D_3+P_3)}{(1+r^E)^3}$$

Are you seeing a pattern?

If we keep this up for an arbitrary number of periods, T, we get that the price of the stock today is

$$P_0 = \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}(D_2)}{(1+r^E)^2} + \frac{\mathbb{E}(D_3)}{(1+r^E)^3} + ... + \frac{\mathbb{E}(D_T+P_T)}{(1+r^E)^T}$$

If we keep this up forever, we get

$$P_0 = \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}(D_2)}{(1+r^E)^2} + \frac{\mathbb{E}(D_3)}{(1+r^E)^3} + ... \tag{8.3}$$

Let's digest what we've just shown. First, the price of a stock today is the sum of the discounted dividends, i.e., cash flows. In other words, our fundamental valuation relation still holds. Second, it doesn't matter for the stock's 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, i.e., dividends. That is, our fundamental valuation relation always holds!

Equation 8.3 is called the **dividend discount model**, which defines the cash flows in our fundamental valuation relation as dividends. We can interpret equation 8.3 as either (i) the value of all the firm's equity or (ii) the value of one share, depending on how we measure dividends. If dividends are measured per share, for example \$0.14 per share, then the value we get is the price per share. If the dividends are measured in total, for example \$52 million, then the value we get is of all the equity, i.e., the market capitalization.

The dividend discount model says that the value of equity is a perpetuity whose cash flows are dividends. At first glance, the model appears pretty naive. In fact, it's fairly rich. We'll look at some of its implications below.

Copyright © 2020 Michael R. Roberts

8.5 Application: The Gordon Growth Model

Unless we have a lot time, the dividend discount model as written in equation 8.3 is impossible to implement in practice because we need to forecast dividends *forever*. To make the model useful, we need to make a simplifying assumption about how dividends evolve over time. The most common assumption is that dividends grow at a constant rate, g. In this case, the dividend discount model says the stock price is a growing perpetuity.

$$\begin{aligned} P_0 &= \frac{\mathbb{E}(D_1)}{(1+r^E)} + \frac{\mathbb{E}(D_1(1+g))}{(1+r^E)^2} + \frac{\mathbb{E}(D_2(1+g)^2)}{(1+r^E)^3} + ... \\ &= \frac{\mathbb{E}(D_1)}{r^E - g} \end{aligned} \tag{8.4}$$

This version of the dividend discount model is called the **Gordon growth model**.

In reality, dividends don't grow at a constant rate. When they increase, they tend to step up by a constant amount each quarter or year. They also decrease and are zero for most firms that don't pay dividends. So, take this assumption for what it is, an assumption that simplifies things but allows for interesting insights.

8.5.1 Expected Returns

Let's solve the Gordon growth model (equation 8.4) for the expected stock return, r^E.

$$r^E = \frac{\mathbb{E}(D_1)}{P_0} + g \tag{8.5}$$

This result says that stock returns are the sum of the **dividend yield**, D_1/P_0, and the dividend growth rate, g. But equation 8.2 above says that the expected stock return equals the dividend yield plus price appreciation. Aha! This means that price appreciation in a stock comes from the growth of future cash flows, i.e., dividends. We'll revisit this result below.

Until then, let's explore equation 8.5 using the data in table 6 for Microsoft.

The table shows the stock price as of February 2, 2022 and the return on equity (ROE) computed as the ratio of the **trailing twelve months** (**ttm**) of earnings to book equity.[3] The forward estimates contain analysts expectations over the next year (or two) and combines

[3]Strictly speaking, **ttm** uses data from the four most recently reported quarters, which in this case includes December 31, 2020 to December 31, 2021.

Copyright © 2020 Michael R. Roberts

	2021A	2022P	2023P
Current information			
Share price ($)	311.21		
TTM Return on equity (ROE)	49.05%		
Forward estimates			
Earnings per share, EPS ($)	8.05	9.35	10.75
Price-to-earnings ratio, P/E	33.28		
Dividend per share, DPS		2.48	
Dividend payout ratio (d)	26.52%		

Table 6: Microsoft Data (Source: Yahoo! Finance, February 2, 2022)

it with current information, e.g., P/E ratio. The "A" and "P" suffixes on the year labels correspond to actual and projected, respectively.

Microsoft's forward dividend yield is the ratio of the forward dividend per share to its current price per share, or $2.48/311.21 = 0.80\%$. To estimate the dividend growth rate, g, we'll take two approaches. First, we can estimate the growth rate with the analysts' earnings forecasts presented in Table 6. The projected growth rates over the next two years are 16.1% and 15.0%, or 15.6% on average. If we assume dividend payouts are a constant fraction of earnings, then dividends will grow at the same rate as earnings.[4] Combining average earnings growth with our earlier estimate implies

$$r^E = \frac{\mathbb{E}(D_1)}{P_0} + g = \frac{2.48}{311.21} + 0.156 = 0.1640. \tag{8.6}$$

In other words, Microsoft's expected stock return based on the Gordon growth model is 16.4% per year. Of course, as prices, expected dividends, and future growth change, so too will the expected return.

Another way to estimate dividend growth, g, is by considering how much of the company's earnings are **plowed back** or **reinvested** in the company, and the rate of return on those investments. Here's the intuition. When a firm generates a $1 of earnings (i.e., net income), it distributes a fraction to shareholders and reinvests the rest. The fraction it distributes is called a **dividend payout ratio** or sometimes just **payout ratio**, which we'll denote by d. For Microsoft, the current dividend payout ratio is 26.52%. Therefore, the fraction of earnings it reinvested in the company is $1 - d$, or 73.48%.

[4]Imagine earnings, E, grow at rate g each year. Then annual earnings are E_1, $E_1(1+g)$, $E_1(1+g)^2$,... If we distribute a constant fraction, d, of earnings each year, then the dividend stream is dE_1, $dE_1(1+g)$, $dE_1(1+g)^2$,... In other words, dividends grow at the same rate, g, as earnings.

Copyright © 2020 Michael R. Roberts

The book return on equity, ROE, measures how much equityholders earn from each dollar they invest. In this case, we measure it as the ratio of the trailing twelve months earnings to the book value of equity from the most recent balance sheet (December 31, 2021). If Microsoft took each dollar of earnings and invested it all back into the company, those earnings, and by extension the dividends, would grow at a rate measured by ROE, or 49.5% per annum according to table 6. But, the payout ratio from table 6 tells us Microsoft only expects to invest \$0.7348 of each dollar in earnings. Therefore, earnings and dividends grow at rate

$$g = (1 - d) \times ROE = (1 - 0.2652) \times 0.495 = 0.3637.$$

Combining this growth rate with our earlier estimates implies

$$r^E = \frac{\mathbb{E}(D_1)}{P_0} + g = \frac{2.48}{311.21} + 0.3637 = 0.3717. \tag{8.7}$$

This growth estimate leads to an estimated expected return for Microsoft of 35.8% per year. That's a big difference from our previous estimate of 16.4% all because we used a different approach to estimate earnings/dividend growth. Which one's correct? Neither. The relevant question is which one's closer to the truth?

It's easier to dismiss the higher estimate as being unreasonably high for at least two reasons. First, the estimated growth rate, 35%, is more than double that of the two most recent years (see table 6). There is no obvious reason, based on an understanding of Microsoft's business and scale, that earnings growth will increase so dramatically going forward. Second, the estimates of the growth rate g are estimates of a constant growth rate *forever*. No firm will have earnings growth of 35% per year forever without taking over the entire global economy. (Of course, no firm will grow at 15.6% per year forever either!) Rather, Microsoft's stellar performance in 2021 has likely resulted in an abnormally high ROE that inflates our growth estimate.

Some clarifying comments are in order on what we just did. First, as expected dividends, prices, and growth rates change, so too will expected returns. Second, be careful interpreting the return estimate, which is Microsoft's *expected* return. Microsoft's *actual* return varies from year to year and will probably never exactly equal 16.4% (or 35.0%). The expected return is just our best guess each year given the information we have.

More generally, equation 8.5 tells us how managers can improve their stock returns, or what investors should look for in a stock. Paying more in dividends and growing these dividends is what drives stock returns. Of course, these two actions are at odds with one

Copyright © 2020 Michael R. Roberts

another. Growth comes from reinvesting in the company. If firms pay out their earnings in dividends instead of reinvesting their earnings in the company, then it will be difficult to grow.

8.5.2 Dividend Growth

Above we noted that the growth rate of dividends and the stock price appreciate were equal. Let's see this more clearly by solving the Gordon growth model (equation 8.4) for the dividend growth rate, g.

$$\begin{aligned} g &= r^E - \frac{\mathbb{E}(D_1)}{P_0} \\ &= \frac{\mathbb{E}(D_1 + P_1) - P_0}{P_0} - \frac{\mathbb{E}(D_1)}{P_0} \\ &= \frac{\mathbb{E}(P_1 - P_0)}{P_0} \end{aligned} \tag{8.8}$$

Equation 8.8 shows that the expected **price appreciation** or **capital gains** of the stock is equal to the growth rate of dividends, g. If a company wants to increase its stock price, then it must increase the growth rate of their dividends, i.e., cash flows to shareholders.

This result does *not* mean simply increasing dividends today, but forever. Increasing future dividend payments requires more earnings from which to pay the dividends, which in turn requires investment to grow the revenues that generate earnings. But, paying out more dividends reduces reinvestment responsible for growth. This is the balancing act corporate managers face - paying out enough of their earnings to shareholders without sacrificing the future growth potential of the company. This is also the challenge facing investors, who need to identify firms with growth potential and the ability to execute this balancing act.

8.5.3 Most Companies Don't Pay Dividends

What about stocks that don't pay dividends? As of 2021, more than half of U.S. publicly traded companies do not pay a dividend. If a company doesn't pay a dividend, how can it have value? 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. Also, liquidating dividends are often paid to shareholders when a company is dissolved, so even for companies that struggle, money is often returned to shareholders.

Copyright © 2020 Michael R. Roberts

The point is that in order for a stock to have value, there must be some hope of getting money from the company in the future. You might be thinking: "I don't need money from the company. I just need money from the person to whom I sell my stock." And, maybe that person is thinking the same thing. But, at some point, this logic is going to unravel because it is tantamount to a ponzi scheme; one in which there is no money at the end of the rainbow, just a hope that you're not the last person to buy. This investment strategy is sometimes referred to as "the greater fool theory," because we have to hope to find someone even more foolish than us to buy a stock that will never distribute any money.

One glaring weakness of the dividend discount model (and its special case the Gordon growth model) is that dividends are not the only way firms can distribute money to shareholders. Share buybacks or repurchases are a popular and important way for companies to get money to their shareholders. This means you could have companies distributing large sums of money to their shareholders via share repurchase programs and because they never pay a dividend have zero value according to the dividend discount model. Our next model addresses this deficiency.

8.6 Total Payout Model

The total payout model recognizes that firms get money to shareholders through dividends *and* share repurchases. According to the total payout model,

$$P_0 = \frac{\mathbb{E}(d_1 E_1)}{(1+r^E)} + \frac{\mathbb{E}(d_2 E_2)}{(1+r^E)^2} + \frac{\mathbb{E}(d_3 E_3)}{(1+r^E)^3} + ... \tag{8.9}$$

where d_t and E_t are the **total payout ratio** and **earnings** at time t. Now, d_t measures the fraction of earnings distributed to shareholders either as a dividend *or* a share repurchase. Earnings refer net income at the bottom of the income statement. These are the profits to which shareholders have a claim. The total payout model measures cash flows to shareholders differently than than the dividend discount model.[5]

Table 7 presents Microsoft's shareholder distributions and issuances, along with their earnings and payout ratios. The column labeled 2022 corresponds to the trailing twelve months as of February 10, 2022. The columns 2019, 2020, and 2021 correspond to June 30 of each year, Microsoft's fiscal year end. The letters in parentheses are abbreviations for each line item, e.g., E = Earnings, I = Issuances, etc.

[5] A 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.

Copyright © 2020 Michael R. Roberts

	2019	2020	2021	2022
Earnings (E)	39,240	44,281	61,271	71,185
Share Issuances (I)	1,142	1,343	1,693	1,749
Share Repurchases (R)	19,453	22,968	27,385	29,224
Dividends (D)	13,811	15,137	16,521	17,293
Total payout (D+R)	33,264	38,105	43,906	46,517
Total net payout (D+R-I)	32,122	36,762	42,213	44,768
Dividend payout ratio (D/E)	35.2%	34.2%	27.0%	24.3%
Total payout ratio ((D+R)/E)	84.8%	86.1%	71.7%	65.3%
Total net Payout ratio ((D+R-I)/E)	81.9%	83.0%	68.9%	62.9%

Table 7: Microsoft Shareholder Cash Flows ($mil)

Microsoft spends 40%-70% more on share repurchases than dividends. We'll postpone the reasons why for later when we discuss corporate financial policy (chapter 11). The point worth emphasizing now is that the their is a big difference between Microsoft's dividend payout and its total payout, or total net payout for that matter. For 2022, Microsoft's dividend payout is 24.3%

Like the dividend discount model, the total payout model requires some assumptions to make it usable. First, we'll assume that the payout ratio is constant over time so $d_t = d$ for all t. Second, we'll assume that earnings grow at a constant rate, g. Under these assumptions, equation 8.9 simplifies to a growing perpetuity (equation 2.8).

$$\begin{aligned} P_0 &= \frac{d\mathbb{E}(E_1)}{(1+r^E)} + \frac{d\mathbb{E}(E_1)(1+g)}{(1+r^E)^2} + \frac{d\mathbb{E}(E_1)(1+g)^2}{(1+r^E)^3} + ... \\ &= \frac{d\mathbb{E}(E_1)}{r^E - g} \end{aligned} \quad (8.10)$$

8.6.1 Growth

If earnings don't grow (or contract), then $E_1 = E_2 = E_3 = ...$, and equation 8.10 becomes

$$P_0 = \frac{d\mathbb{E}(E_1)}{r^E}.$$

The right side of this equation is the present value of a perpetuity in which the cash flow is $d\mathbb{E}(E_1)$ and the discount rate is r^E (see equation 2.8). Without earnings growth, and assuming the expected equity return is constant, the price of a stock is always the same. For

Copyright © 2020 Michael R. Roberts

example, the price of a stock six years from today is

$$P_6 = \frac{d\mathbb{E}(E_7)}{r^E} = \frac{d\mathbb{E}(E_1)}{r^E},$$

where the second equality follows from the lack of earnings growth (i.e., all earnings are the same).

This is why managers are so focused on revenue growth. It's the only way they can consistently increase the price of their stock. Yes, they can cut costs to boost earnings in the short-run, but cost-cutting can only last for so long. Eventually, firms have to grow if they want their value to grow.

8.7 Application: The P/E and PEG Ratios

We can solve equation 8.10 for the **forward** or **prospective price-to-earnings ratio** (i.e., **P/E ratio**).

$$\frac{P_0}{\mathbb{E}(E_1)} = \frac{d}{r^E - g} \tag{8.11}$$

The P/E ratio measures how much investors are willing to pay for $1 of next year's expected earnings, also known as **forward** or **prospective** earnings . (In practice, people sometimes compute the P/E ratio using the current year's earnings to avoid having to come up with an earnings forecast.)

Equation 8.11 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. Investors like receiving more cash flow and therefore pay more today.

2. The cost of equity capital (r^E). As the expected equity return increases, the P/E ratio decreases. 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. Investors like growth which leads to higher future prices and therefore pay more today.

There are several ways to measure the P/E ratio which differ by how we measure earnings. A simple method is to take the current market capitalization and divide it by the earnings over the most recent year (i.e., trailing twelve months). Similarly, we could take the current price per share and divide it by the trailing twelve months **earnings per share** (**EPS**).

Copyright © 2020 Michael R. Roberts

Because the earnings in equation 8.11 are next year's earnings, we might instead use a forward looking measure of earnings, such as next year's forecasted earnings or earnings per share.

Let's examine's Microsft's trailing twelve month P/E ratio, in the context of its peers, the broader stock market, and its history. Table 8 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.

	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 8: Trailing Twelve Month P/E Ratios

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 8.11, this increase is driven by a larger payout ratio, greater earnings growth, lower expected return or some combination. As the payout ratio has been declining over time, and is expected to continue next year, (see table 7) chances are the increasing P/E ratio is due to the increased expectations for earnings growth and, perhaps, lower expected returns.

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 that Amazon is a "better" company than IBM? Not necessarily, though Amazon's recent performance has been much better than IBM's. Equation 8.11 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? Unfortunately, it is difficult to answer that question because different P/E ratios could just be masking different growth expectations, expected returns, and future payout ratios.

Copyright © 2020 Michael R. Roberts

A few things to keep in mind. First, negative P/E ratios are meaningless; investors don't pay to lose money. So, if earnings are negative, the P/E ratio is undefined. Second, be careful of interpreting very large P/E ratios (e.g., 500); they could just reflect very low current or short-term earnings; dividing by a small number can lead to a very large number.

8.7.1 What About Growth?

When we consider the three determinants of P/E ratios - d, r^E, and g from equation 8.11 - the latter two are responsible for most of the variation across firms. In other words, payout ratios don't vary much across firms when compared to the variation in 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.

$$\text{PEG Ratio} = \frac{\text{P/E Ratio}}{\text{Earnings Growth Rate (measured in \%)}} \tag{8.12}$$

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.

$$\text{MSFT PEG Ratio} = \frac{32.3}{15} = 2.15$$
$$\text{AMZN PEG Ratio} = \frac{59.9}{48.8} = 1.23$$

Loosely speaking, the PEG ratio is telling us the price stock market participants are willing to pay $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

Copyright © 2020 Michael R. Roberts

is because of the greater earnings growth. Relative to Microsoft, the PEG ratio suggests Amazon is cheap.

A popular rule of thumb uses 1.0 as a benchmark. Companies with PEG ratios above one are overvalued, below one are undervalued. While convenient, this rule is ad hoc, and there is no compelling empirical evidence suggesting that the PEG ratio can be used to forecast future stock returns.

8.8 Flow to Equity (a.k.a., DCF)

The last equity valuation model we'll discuss bears a lot of similarities to the discounted cash flow analysis we performed for capital budgeting (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**).

$$Value_0 = \frac{\mathbb{E}(FCFE_1)}{1+r^E} + \frac{\mathbb{E}(FCFE_2)}{(1+r^E)^2} + \frac{\mathbb{E}(FCFE_3)}{(1+r^E)^3} + ... \tag{8.13}$$

Levered free cash flow is defined as

$$\begin{aligned} FCFE &= \overbrace{(Sales - Expenses - D\&A) \times (1-\tau) + D\&A - NLTI - NWCI}^{\text{Unlevered free cash flow } (FCF)} \\ &\quad - \underbrace{(1-\tau) \times Interest + NDI}_{\text{Cash flow to creditors}}, \end{aligned} \tag{8.14}$$

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). So, levered free cash flows starts with unlevered free cash flow, subtracts any payments to creditors - interest and principal repayment - and adds any receipts from creditors - new borrowing. Because interest is tax deductible, we subtract off the after-tax interest expense.

Just like our first two stock models, the flow to equity method runs into the problem of an infinite number of cash flows. The solution is to assume that all cash flows beyond a certain time period, T (e.g., five or seven years), grow at a constant rate thereafter. Figure 8.2 illustrates this assumption.

Leverage free csh flows can change in an arbitrary manner over the short-term, i.e., up to and including period T. However, from period $T+1$ onward, we assume the cash flows

Copyright © 2020 Michael R. Roberts

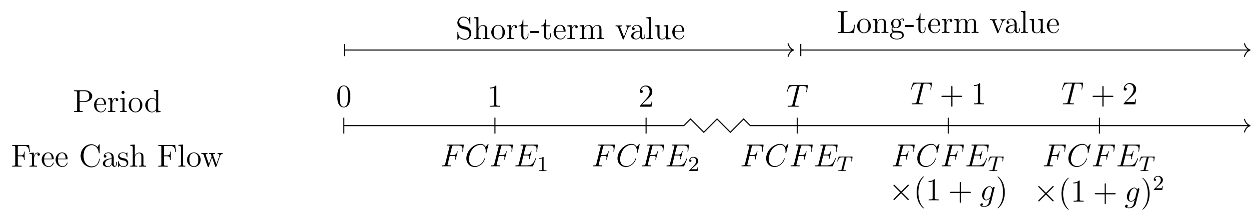

Figure 8.2: Flow to Equity Method

grow at a constant rate g. Mathematically, equity value can be written as

$$\begin{aligned} Value_0 &= \overbrace{\frac{FCFE_1}{(1+r^E)} + \frac{FCFE_2}{(1+r^E)^2} + ... + \frac{FCFE_T}{(1+r^E)^T}}^{\text{Short-term value}} \\ &+ \underbrace{\frac{FCF_T(1+g)}{(1+r^E)^{T+1}} + \frac{FCFE_T(1+g)^2}{(1+r^E)^{T+2}} + ...}_{\text{Long-term value}}. \end{aligned} \tag{8.15}$$

The constant growth after year T implies the cash flows from period $T+1$ onward are a growing perpetuity. From equation 2.8 in chapter 2, the value of equity as of period T is

$$\begin{aligned} Value_T &= \frac{\overbrace{FCFE_T(1+g)}^{FCFE_{T+1}}}{1+r^E} + \frac{\overbrace{FCFE_T(1+g)^2}^{FCFE_{T+2}}}{(1+r^E)^2} + \frac{\overbrace{FCFE_T(1+g)^3}^{FCFE_{T+3}}}{(1+r^E)^3} + ... \\ &= \frac{FCFE_T(1+g)}{r^E - g}. \end{aligned} \tag{8.16}$$

The timing can be confusing so let's explain. We are computing the value of the equity as of period T. This value is a function of all the *future* cash flows from $T+1$ onward, which happen to be growing at a constant rate g. The $T+1$ cash flow is $FCF_T(1+g)$, which is our numerator in equation 8.16. The denominator is the discount rate, r^E, less the cash flow growth rate, g.

Let's combine the result in equation 8.16 with equation 8.15.

$$Value_0 = \underbrace{\frac{FCFE_1}{(1+r^E)} + \frac{FCFE_2}{(1+r^E)^2} + ... + \frac{FCFE_T}{(1+r^E)^T}}_{\text{Short-term value}} + \underbrace{\frac{FCFE_T(1+g)}{r^E - g}}_{\text{Long-term value, } Value_T} \tag{8.17}$$

The last term in equation 8.17, $Value_T$, is called the **terminal value** or **continuation value**. It represents the value of the equity at date T, which is itself the sum of all future discounted free cash flows to shareholders.

Copyright © 2020 Michael R. Roberts

8.8.1 Terminal Value Growth and the Short-term Horizon

What should T be in practice? That is, how far out should we try to forecast free cash flows before throwing in the towel and assuming a constant growth rate? 1 year? 5 years? 10 years? 100 years? There is no right answer and, importantly, your choice doesn't matter from a theoretical perspective. However, a useful guide is that T should be chosen so that the growth of the cash flows has settled into a **long-run equilibrium** or **steady-state**, something that is not changing much. That's a mouthful so let's elaborate.

The terminal value, $Value_T$, assumes that future cash flows grow at a constant rate, g, *forever*. The "forever" is important because it constrains what number we can use for g. Specifically, g shouldn't be (much) larger than the growth rate of the economy, as measured by **Gross Domestic Product** or **GDP**. For example, U.S. GDP growth has been about 3% per year.

Why this limit? If we use a number larger than GDP growth, then our firm will grow to eventually become the entire economy. This outcome is silly so we constrain the choice of g to be approximately equal to the growth rate of the economy, if not less or even negative in say a contracting industry (e.g., print media at the time of this writing).

Given this constraint on g, we want to choose a short-term forecast horizong, 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 9.

	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 9: Example Cash Flow Growth Profiles

Table 9 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

Copyright © 2020 Michael R. Roberts

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, 10 to 20 years, makes more sense before relying on a terminal value.

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.

8.9 Historical Data and Risk-Reward Estimates

8.9.1 Average Returns

The models we discussed provide different ways to value a stock and, by extension, estimate a stock's expected return, r^E. But, rather than relying on a particular model to estimate the expected return, we can just use historical data. Statistics tells us that the **arithmetic average**, or just **average**, provides a pretty good estimate of what one might expect to happen. In other words, we can estimate a stock's expected return looking at its average historical return.

Remember that 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 can be expressed as

$$\text{Average} = \frac{r_1 + r_2 + ... + r_N}{N}, \tag{8.18}$$

where $r_1, r_2, ..., r_N$ are the N historical returns in our sample.

For example, consider estimating the expected return for Microsoft. Table 10 shows a partial view of Microsoft's historical monthly returns since it went **IPO**, i.e., became a public company through an **Initial Public Offering**.

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 estimate is that it is not very precise. As we'll see in a moment, stock returns vary a lot over

Copyright © 2020 Michael R. Roberts

Date	Return
19860430	0.172727
19860530	0.077519
19860630	-0.115108
⋮	⋮
20211029	0.176291
20211130	-0.001236
20211231	0.017333

Table 10: Microsoft Monthly Historical Returns

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 **initial public offering** or **IPO**. This approach gets us a sample with 431 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.

8.9.2 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 in this book. For now, we'll focus on perhaps the most intuitive and common measure, **standard deviation** or

Copyright © 2020 Michael R. Roberts

volatility. The volatility of a stock's return measures how much those returns are expected to vary around its 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.

$$\text{Standard deviation} = \sqrt{\underbrace{\frac{1}{N-1}\left[(r_1-\bar{r})^2+(r_2-\bar{r})^2+...+(r_N-\bar{r})^2\right]}_{\text{Variance}}} \tag{8.19}$$

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**, indicated in equation 8.19. We focus on standard deviations because the units are easier to interpret. (Variance is in units squared.)

The equation looks a little scary but it is quite 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 8.18.
2. Subtract the average return from each return in the sample. This gives us **deviations from the average**: $(r_1-\bar{r}),(r_2-\bar{r}),...,(r_N-\bar{r})$.
3. Square each of these deviations from the average: $(r_1-\bar{r})^2,(r_2-\bar{r})^2,...,(r_N-\bar{r})^2$.
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.

Using monthly returns from 2017 to 2021, the monthly volatility of Microsoft is 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. This tells us that the stock returns can make very large swings.

Figure 8.3 presents the time series of monthly returns for Microsoft over the period 2017 to 2021. The dashed red line indicates the average return over this period, 3.11%. The figure illustrates just how much monthly stock returns bounce around over time from a low

Copyright © 2020 Michael R. Roberts

of -8.41% in December 2018 to a high of 17.63% in October 2021. The bouncing around is measured by the 5.24% monthly volatility.

Figure 8.4 presents a histogram of Microsoft's monthly stock returns and another way of viewing just how volatile stock returns are. 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 quite spread out around that average, reflecting the large volatility in returns.

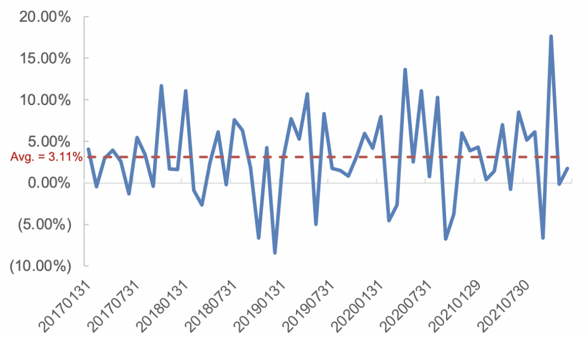

Figure 8.3: Microsoft Monthly Stock Returns

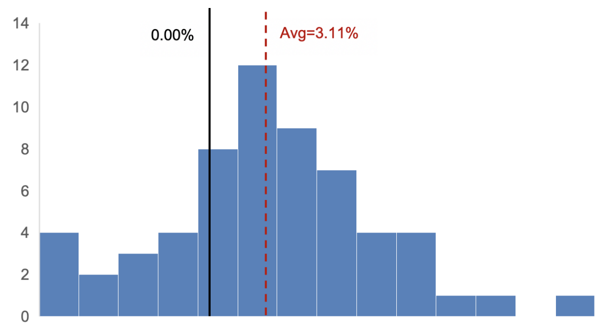

Figure 8.4: Microsoft Monthly Stock Return Histogram

To summarize, one way to estimate expected returns without a specific model 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 difficulty because returns are highly volatile; i.e., they have a large standard deviation relative to their average.

8.9.3 Annualizing Estimates

Our expected return and volatility estimates reflect *monthly* returns because that is the frequency of the data with which we are working. 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.[6] Specifically,

$$\begin{aligned}\text{Average annual return} &= 12 \times \text{Average monthly return} \\ \text{Average annual volatility} &= \sqrt{12} \times \text{Average monthly volatility}\end{aligned}$$

[6]This 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.

Copyright © 2020 Michael R. Roberts

The "12" and "$\sqrt{12}$" in the equations above are the scale factors for converting the monthly average return and volatility into annual average return and volatility, respectively. 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, when we want to convert an estimated periodic expected return and volatility, we scale by the number of periods in a year, k.

$$\text{Average annual return} = k \times \text{Average periodic return} \tag{8.20}$$

$$\text{Average annual volatility} = \sqrt{k} \times \text{Average periodic volatility} \tag{8.21}$$

For example, when working with weekly returns, $k = 52$. When working with daily returns, $k = 252$. (There are only 252 **trading days** in a calendar year because stock markets are closed on weekends and a number of federal holidays.)

You might be wondering why we don't compound our periodic return to get an annual return like so.

$$\text{Average annual return} = (1 + \text{Average periodic return})^k - 1$$

When we estimate *expected returns*, we want to use arithmetic averages (equation 8.18) and this scaling makes sense. When we are trying to understand *historical 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.

8.10 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 shareholders ownership stake is proportional to the number of shares they own.
- 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,

Copyright © 2020 Michael R. Roberts

2. Total payout model, and
3. Flow to equity model.

- 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/deprecation is in turn equal to dividend or earnings growth according to the dividend discount and total payout models, respectively. Simply put, for stock prices to increase, dividends or earnings must grow.

- Using the total payout model, we saw that stock P/E ratios are a positive function of how much the firm distributes to shareholders and the growth of earnings, and a negative function of expected returns. 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 with 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.

- 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.

8.11 Problems

8.1 The table below presents market data for Goldman Sachs (GS). The Price column presents the closing price, i.e., the price of the day's last trade.

Date	Dividend ($/share)	Price ($/share)
8/30/2021	0	413.60
8/31/2021	2	413.51
9/1/2021	0	413.66

Using this data, answer the following questions.

a. What are the one-day price returns?

Copyright © 2020 Michael R. Roberts

b. What are the one-day dividend yields?

c. What are the one-day total returns?

d. 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?

e. 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?

8.2 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.

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 aritmetic 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 short-sold 100 shares of GameStop stock at the closing price on January 22nd, how much money would you have received ignoring any transaction costs?

Copyright © 2020 Michael R. Roberts

If you were forced to buy and return the stock on January 29th at the closing price, perhaps because of short squeeze, how much money would you have to spend? What would your net gain/loss be from this transaction?

8.3 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
20201230	705.67
20210129	793.53
20210226	675.50
20210331	667.93
20210430	709.44
20210528	625.22
20210630	679.70
20210730	687.20
20210831	735.72
20210930	775.48
20211029	1,114.00
20211130	1,144.76
20211231	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.?

8.4 The table below presents daily stock market data for Amazon (AMZN).

Copyright © 2020 Michael R. Roberts

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.

a. What was the split ratio for the stock split Amazon executed on June 6th?

b. What was the June 6th return on Amazon stock?

c. Did Amazon shareholders make or lose money on June 6th?

8.5 Short-term v long-term investors.

8.6 Selling stocks and capital gains taxes

8.7 Dividend taxes with stocks

8.8 As of October 7, 2022, Yahoo! Finance shows the following financial information for Bank of America.

Price per share ($)	31.46
Forward annual dividend per share ($)	0.88

Copyright © 2020 Michael R. Roberts

Chapter 9

Investing: Portfolios

Fundamental value relation

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

A portfolio is a collection of assets - e.g., stocks, bonds, real estate, commodities, projects? What happens to our returns - realized and expected - and risk when we own a portfolio of assets?

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* the sum of the risk of the assets in the portfolio,

Copyright © 2020 Michael R. Roberts

- 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 spreadsheet program. The key is not to lose sight of the intuition, which we'll emphasize throughout.

9.1 Portfolio Returns

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 (ticker symbol MSFT) and Ferrari (ticker symbol RACE). Specifically, the table shows the end of month share price, dividend per share, and monthly total *realized return* for the start and end of the five year period encompassing 2017 through 2021.

	Microsoft (MSFT)			Ferrari (RACE)		
Date	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

We've seen how to compute returns for bonds (chapter 7) and stocks (chapter 8) . More

Copyright © 2020 Michael R. Roberts

generally, the return on any asset over the period $t-1$ to t can be expressed as follows.

$$r_{t-1,t} = \frac{V_t - V_{t-1} + CF_t}{V_{t-1}} \tag{9.1}$$

Equation 9.1 says that the return on an investment over some time period, $t-1$ to t, is equal to the value of the investment at time t (V_t) minus the value of the investment at time $t-1$ (V_{t-1}) plus any cash flow the investment produces during the period (CF_t) all divided by the start of period investment value (V_{t-1}).

For stocks and bonds, the values, V_t and V_{t-1}, are the prices. The cash flows for bonds are the periodic coupons (if any) and principal payment. The cash flows for a stock are the dividends, if any.

Let's compute the return to a portfolio consisting of both Microsoft and Ferrari during November 2021. To do so, we have to know how much money is invested in each stock, which in turn requires knowing how many shares of each stock are in our portfolio and the price per share. Let's assume we have five shares of Microsoft and ten shares of Ferrari.

The value of our portfolio at the end of October 2021 (i.e., the start of the period) is

$$V_{t-1} = \overbrace{5}^{\text{MSFT shares}} \times \underbrace{331.62}_{\text{MSFT price per share}} + \overbrace{10}^{\text{RACE shares}} \times \underbrace{237.17}_{\text{RACE price per share}} = \$4,029.80.$$

The value at the end of November (i.e., the end of the period) is

$$V_t = 5 \times 330.59 + 10 \times 260.46 = \$4,257.55.$$

Note the number of shares didn't change over the month because we've assumed that we don't buy or sell any stock. Only the prices of Microsoft and Ferrari stocks have changed.

The cash flow during the period is the dividend per share that Microsoft paid times the number of Microsoft shares in our portfolio.

$$CF_t = 5 \times 0.62 = \$3.10$$

So, the return on our portfolio for November 2021 is

$$r_{t-1,t} = \frac{4,257.55 - 4,029.80 + 3.10}{4,029.80} = 0.0573,$$

or 5.73%.

An equivalent and more convenient way to compute portfolio returns recognizes that the return on a portfolio is just a weighted average of the returns to each asset in the portfolio.

Copyright © 2020 Michael R. Roberts

The weights are the relative amount invested in each asset. That's a mouthful so let's unpack it slowly.

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

$$\begin{aligned} w_{t-1}^{MSFT} &= \frac{5 \times 331.62}{5 \times 331.62 + 10 \times 237.17} = \frac{1,658.10}{4,029.80} = 0.4115, \text{ and} \\ w_{t-1}^{RACE} &= \frac{10 \times 237.17}{5 \times 331.62 + 10 \times 237.17} = \frac{2,371.70}{4,029.80} = 0.5885. \end{aligned}$$

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. Notice the weights sum to one. This is true of every portfolio; **portfolio weights always sum to one**!

The portfolio return is the sum of the weighted returns to each asset, where the weights are the portfolio weights.

$$\begin{aligned} r_{t-1,t} &= w_{t-1}^{MSFT} \times r_{t,t-1}^{MSFT} + w_{t-1}^{RACE} \times r_{t,t-1}^{RACE} \\ &= 0.4115 \times (-0.0012) + 0.5885 \times 0.0982 \\ &= 0.0573 \end{aligned}$$

We get the same portfolio return, 5.73%, as we did above.

More generally, the return on a portfolio during the period $t-1$ to t can be computed as follows.

$$r_{t,t-1} = w_{t-1}^{1} r_{t-1,t}^{1} + w_{t-1}^{2} r_{t-1,t}^{2} + ... + w_{t-1}^{N} r_{t-1,t}^{N} \tag{9.2}$$

The portfolio weights, $w_{t-1}^{1}, ..., w_{t-1}^{N}$, represent the fraction of the portfolio value invested in each of the portfolios N assets as of the start of the period, $t-1$. The asset returns, $r_{t-1,t}^{1}, ..., r_{t-1,t}^{N}$, are the individual asset returns over the period $t-1$ to t. Thus, the portfolio return is a weighted sum of the returns to each asset in the portfolio, where the weights are the fraction of the portfolio value invested in each asset.

9.1.1 Changing the Portfolio Weights

What happens if 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 return? Figure 9.1 presents the different portfolio returns as we change Microsoft's portfolio weight (w_{t-1}^{MSFT}),

Copyright © 2020 Michael R. Roberts

i.e., the fraction of our investment allocated to Microsoft as opposed to Ferrari. (The portfolio weight for Ferrari is simply one minus the Microsoft weight since the weights must always sum to one.)

Specifically, the figure plots the following function,

$$\begin{aligned} r^P &= w_{t-1}^{MSFT} r_{t-1,t}^{MSFT} + \underbrace{\left(1 - w_{t-1}^{MSFT}\right)}_{w_{t-1}^{RACE}} r_{t-1,t}^{RACE} \\ &= \underbrace{r_{t-1,t}^{RACE}}_{y\text{-intercept}} + \underbrace{\left(r_{t,t-1}^{MSFT} - r_{t,t-1}^{RACE}\right)}_{\text{Slope}} w_{t-1}^{MSFT}, \end{aligned}$$

where r^P is the portfolio return.

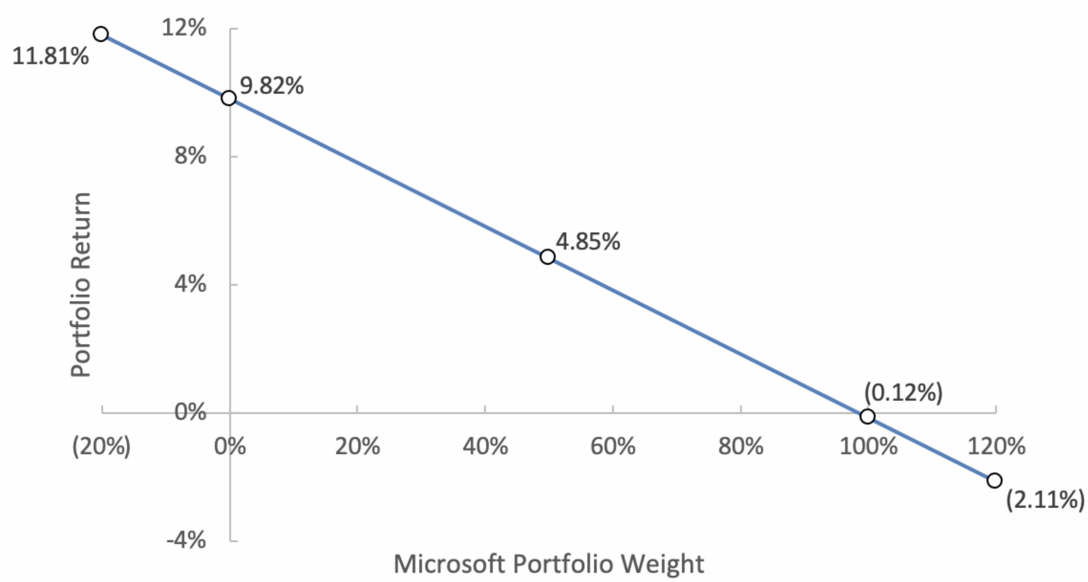

Figure 9.1: Portfolio Returns with Different Portfolio Weights

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 is just equal to 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 ($w_{t-1}^{MSFT} = 1$) so that the portfolio 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.

Copyright © 2020 Michael R. Roberts

9.1.2 Short Positions

You might be wondering: What about the points where the weights on Microsoft are negative or greater than 100%? How can that be? These negative portfolio weights correspond to **short positions**. 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** on Microsoft (i.e., negative, pessimistic, expecting to depreciate in value). 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. 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%. Similarly, if we short-sell $200 worth of Ferrari to buy $1,200 worth of Microsoft, our portfolio return is:

$$\begin{aligned} r_{t-1,t} &= w_{t-1}^{MSFT} \times r_{t,t-1}^{MSFT} + w_{t-1}^{RACE} \times r_{t,t-1}^{RACE} \\ &= \frac{1,200}{1,000} \times (-0.0012) - \frac{200}{1,000} \times 0.0982 \\ &= 1.20 \times (-0.0012) - 0.2 \times 0.0982 \\ &= -0.0211, \end{aligned}$$

or -2.11%.

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 the we typically know the interest rate on a loan. When we short-sell, we don't know the

Copyright © 2020 Michael R. Roberts

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 we 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 possible returns but also risk.

9.1.3 Application: 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 - market cap - of each company in the portfolio scaled by the sum of those market capitalizations.

Let's use our Microsoft-Ferrari portfolio as an example. Recall that the market cap 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.

$$
\begin{aligned}
w_{t-1}^{MSFT} &= \frac{2,489,796,327,600}{2,489,796,327,600 + 43,747,666,690} = 0.9827 \\
w_{t-1}^{RACE} &= \frac{43,747,666,690}{2,489,796,327,600 + 43,747,666,690} = 0.0173
\end{aligned}
$$

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

Copyright © 2020 Michael R. Roberts

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.

9.1.4 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-weighed 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.

Copyright © 2020 Michael R. Roberts

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.

9.1.5 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 expected returns in the place of actual returns in equation 9.2.

$$\mathbb{E}(r) = w^1\mathbb{E}(r^1) + w^2\mathbb{E}(r^2) + ... + w^N\mathbb{E}(r^N) \tag{9.3}$$

Let's use our Microsoft and Ferrari portfolio to illustrate.

The expected return to this two-stock portfolio is

$$\mathbb{E}(r) = w^{MSFT}\mathbb{E}(r^{MSFT}) + w^{RACE}\mathbb{E}(r^{RACE})$$

The average monthly returns to Microsoft and Ferrari between 2017 and 2021 are 3.11% and 2.87%, respectively. We can use these estimates in place of the expected returns to get

$$\mathbb{E}(r) = w^{MSFT}3.11\% + w^{RACE}2.87\%.$$

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.

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

Copyright © 2020 Michael R. Roberts

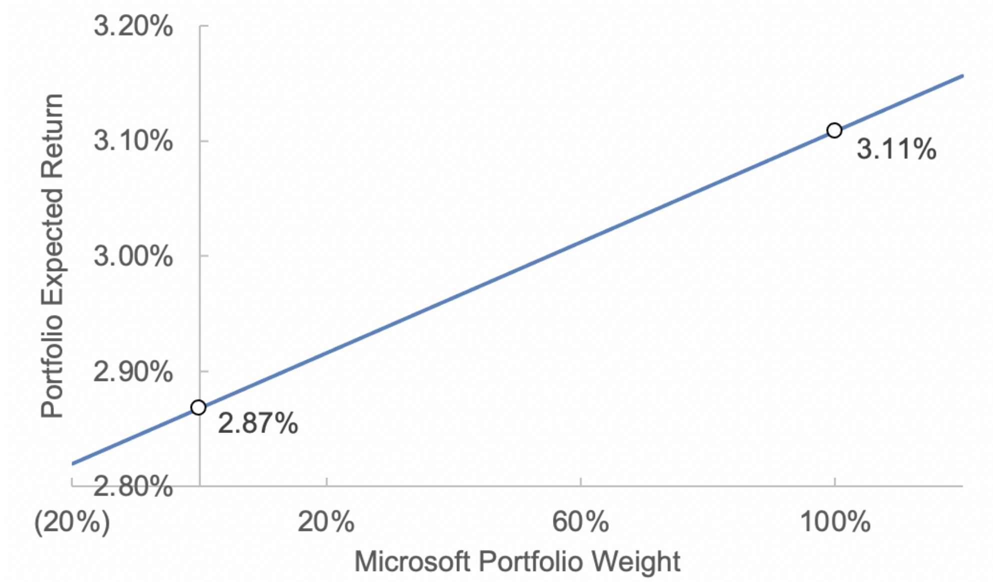

Figure 9.2: Microsoft-Ferrari Portfolio Expected Returns

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 expected returns. Microsoft has a higher expected return and so investing relatively more money in Microsoft than Ferrari (i.e., a higher portfolio weight on Microsoft) leads to higher expected portfolio returns. Of course, the realized future performance of any portfolio is unknown. The expected return is just our best guess.

To bring this discussion back to our fundamental valuation 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

Copyright © 2020 Michael R. Roberts

portfolio can be expressed as follows.

$$\sqrt{(w^{MSFT}SD^{MSFT})^2 + (w^{RACE}SD^{RACE})^2 + 2w^{MSFT}w^{RACE}SD^{MSFT}SD^{RACE}Corr^{MSFT,RACE}} \quad (9.4)$$

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 $Corr^{MSFT,RACE}$, which represents the **correlation** of Microsoft's and Ferrari's returns. Correlation measures the degree to which the two returns are related. Specifically, correlation helps us understand 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**.

Correlations are *always* between -1 and 1. When the correlation is either -1 or 1, then knowing one stock's return implies we can determine exactly 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**. When the correlation is 0, then knowing one stock's return provides no information about the other stock's return. In this case, we say the stock returns are **uncorrelated**.

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.[1] 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%.

I would be remiss if I didn't share an alternative formulation for the volatility of our two asset portfolio that is equivalent to equation 9.4 because it often appears in practice and other books.

$$\sqrt{(w^{MSFT})^2 Var^{MSFT} + (w^{MSFT})^2 Var^{RACE} + 2w^{MSFT}w^{RACE}Cov^{MSFT,RACE}} \quad (9.5)$$

What's different? We recognized that the standard deviation squared is equal to the variance, Var.

$$(SD^{MSFT})^2 = Var^{MSFT} \text{ and } (SD^{RACE})^2 = Var^{RACE} \quad (9.6)$$

[1]Correlation is computed...TODO

Copyright © 2020 Michael R. Roberts

We also recognized the following relation,

$$Corr^{MSFT,RACE} = \frac{Cov^{MSFT,RACE}}{SD^{MSFT}SD^{RACE}}, \tag{9.7}$$

which implies

$$Cov^{MSFT,RACE} = SD^{MSFT}SD^{RACE}Corr^{MSFT,RACE}$$

and where Cov stands for **covariance**. Covariance measures the strength of association between two random variables much like correlation. However ,we tend to focus on correlations in practice because they are easier to interpret being restricted to the interval -1 to 1.[2]

Let's use equation 9.5 to compute the volatility of our 30%-70% Microsoft-Ferrari portfolio just to confirm that we get the same result.

$$\sqrt{(0.3)^2 \times (0.0524)^2 + (0.7)^2 \times (0.0780)^2 + 2 \times 0.3 \times 0.7 \times 0.00263} = 0.0631$$

Figure 9.3 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 is 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. What this picture is showing us is 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**!!!! What's happening is that because the two stocks are only moderately correlated, when Microsoft performs poorly, Ferrari will tend

[2] Covariance units are the units of the variables squared. For our returns, that means percent square, which does not have an intuitive interpretation. In contrast, correlations being restricted to values between -1 and 1 are more easily used to compare different relations.

Copyright © 2020 Michael R. Roberts

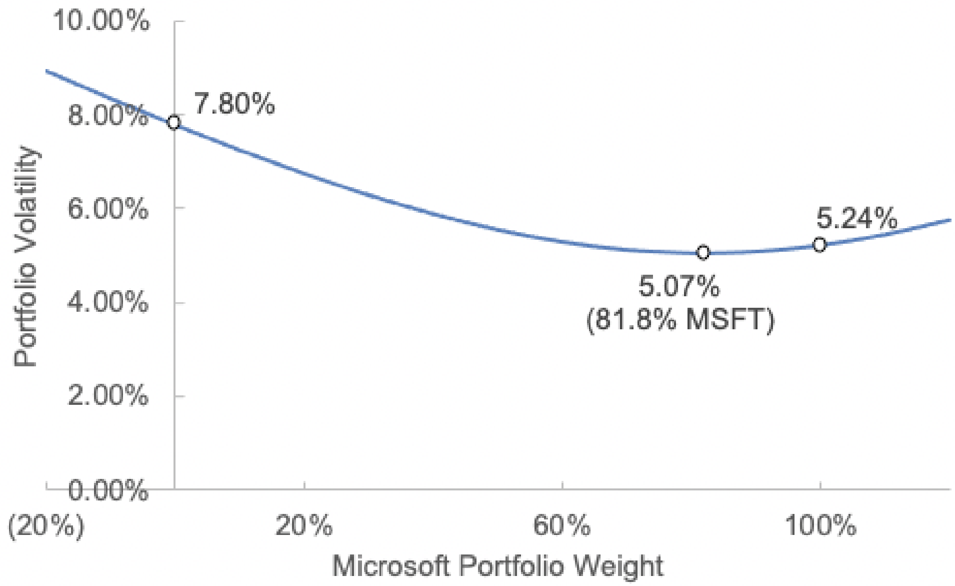

Figure 9.3: Microsoft-Ferrari Portfolio Volatility

to perform poorly but not all the time (and vice versa). In other words, each stock acts as a **hedge**, albeit a modest one, for the other. Losses on one stock are occasionally offset by gains on the other thereby reducing the overall risk of the portfolio.

9.3 Application: Mean Variance Frontier

Let's continue exploring how the expected returns and volatility of a two asset portfolio behave. To shake things up a little and help reinforce concepts, I'm going to focus on Microsoft and Walmart (ticker symbol WMT) during the period 2012-2016. Table 2 displays each stock's summary statistics needed to analyze the behavior of portfolios of these stocks.

	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

I've presented both monthly and annual statistics. The latter are computed using the scaling results (equation 8.21) from chapter 8. As a reminder, we multiply the average return

Copyright © 2020 Michael R. Roberts

by 12 and the standard deviation by $\sqrt{12}$ because there are 12 months in a year. Notice that 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).

Above we saw how expected returns and volatility changed as we changed our portfolio (i.e., changed the portfolio weights). Now let's see both - expected return and volatility - on the same graph to see how the risk-return tradeoff varies with changes in our portfolio. Figure 9.4 presents the **mean-variance frontier** for portfolios of Microsoft and Walmart.[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%.

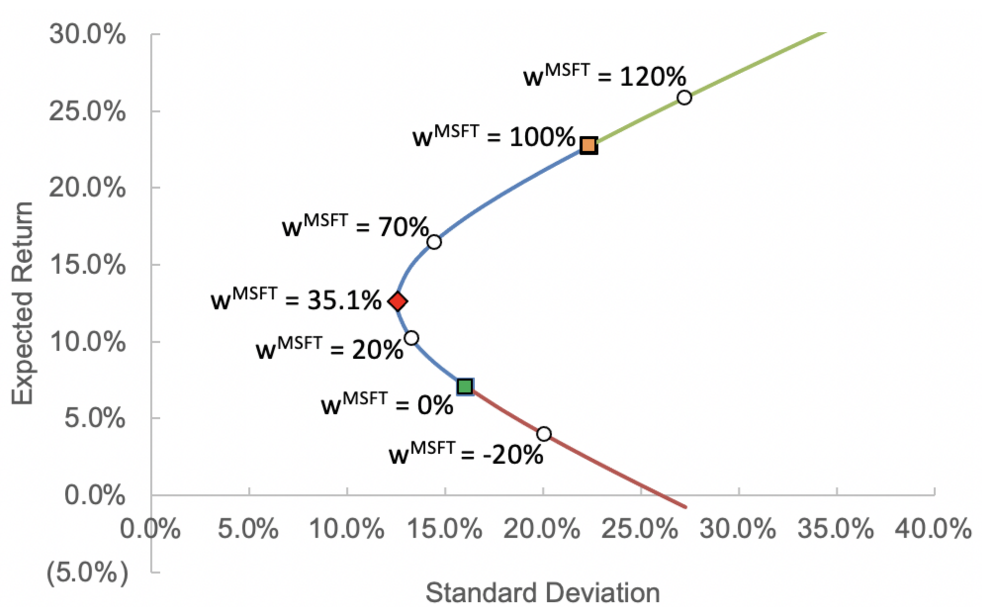

Figure 9.4: Microsoft-Walmart Portfolio Volatility

There are three colored sections of the curve.

[3]This 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.

Copyright © 2020 Michael R. Roberts

1. *Short Microsoft, levered long Walmart.* The orange, lower third of the curve corresponds 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 third 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 third 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.

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 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%.

9.3.1 Efficient and Inefficient Portfolios

This ability to increase expected return and decrease risk suggests another way to look at out portfolios, which is illustrated in Figure 9.5. The figure shows the same curve as in figure 9.4 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

Copyright © 2020 Michael R. Roberts

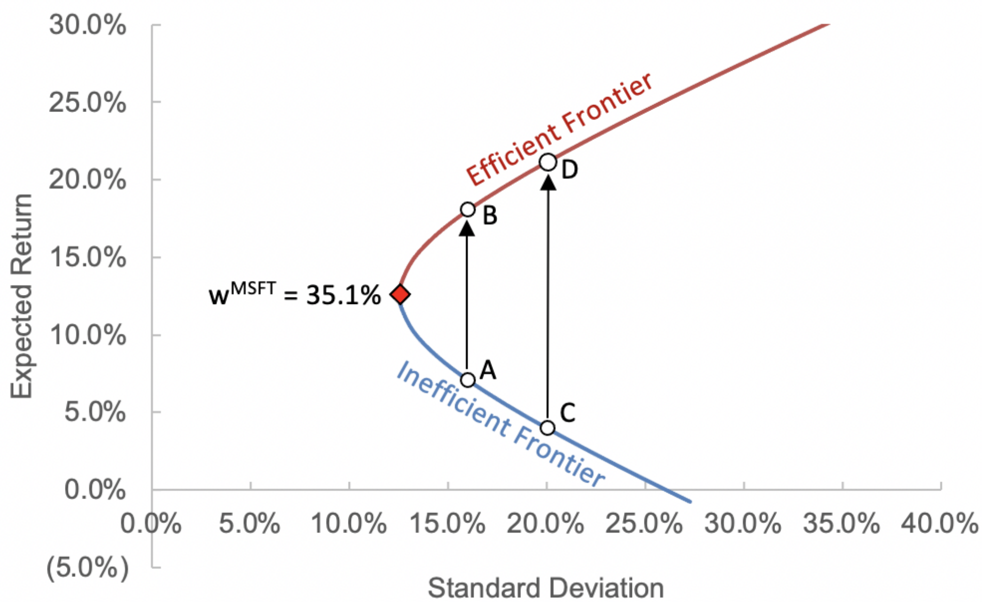

Figure 9.5: Microsoft-Walmart Mean Variance Frontier

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 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 reprsents 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. 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.

[4]There is an exception. The **risk-loving** investor might prefer portfolios A and C to B and D, respectively. However, most people are **risk-averse**, in which case our conclusion makes sense.

Copyright © 2020 Michael R. Roberts

9.3.2 Minimum Variance Portfolio (MVP)

The red diamond in figures 9.4 and 9.5 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.

$$w^A = \frac{(SD^B)^2 - SD^A SD^B corr^{A,B}}{(SD^A)^2 + (SD^B)^2 - 2SD^A SD^B Corr^{A,B}} \tag{9.8}$$

The weight on asset B is simply $1 - w^A$. For completeness, equation 9.8 can be expressed using variances and covariances like so

$$w^A = \frac{Var^B - Cov^{A,B}}{Var^A + Var^B - 2Cov^{A,B}}. \tag{9.9}$$

Let's use equation 9.8 to verify the MVP for our Microsoft-Walmart portfolios.

$$w^{MSFT} = \frac{0.1605^2 - 0.2236 \times 0.1605 \times 0.0733}{0.2236^2 + 0.1605^2 - 2 \times 0.2236 \times 0.1605 \times 0.07333} = 0.3505$$

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 9.5).[6] For the Microsoft-Walmart MVP, the volatility is

$$\sqrt{\frac{.2236^2 \times 0.1605^2 - (0.00263)^2}{0.2236^2 + 0.1605^2 - 2 \times 0.00263}} = 0.1257.$$

Aside from being the lowest risk portfolio, the MVP as seen above in figures 9.4 and 9.5 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.

[5]Since variance is just the square of the standard deviation, the portfolio with the smallest variance is also the portfolio with the smallest standard deviation.

[6]Alternatively, the volatility of the MVP can be computed directly from the variances and covariances of the assets themselves.

$$\sqrt{\frac{Var^A Var^B - (Cov^{A,B})^2}{Var^A + Var^B - 2Cov^{A,B}}}$$

Copyright © 2020 Michael R. Roberts

9.3.3 Changing Correlation

Now let's explore what happens to the mean-variance frontier as the correlation between our two stocks' returns changes. Figure 9.6 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 I know either Microsoft or Walmart's return, I can compute the other stock's 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.

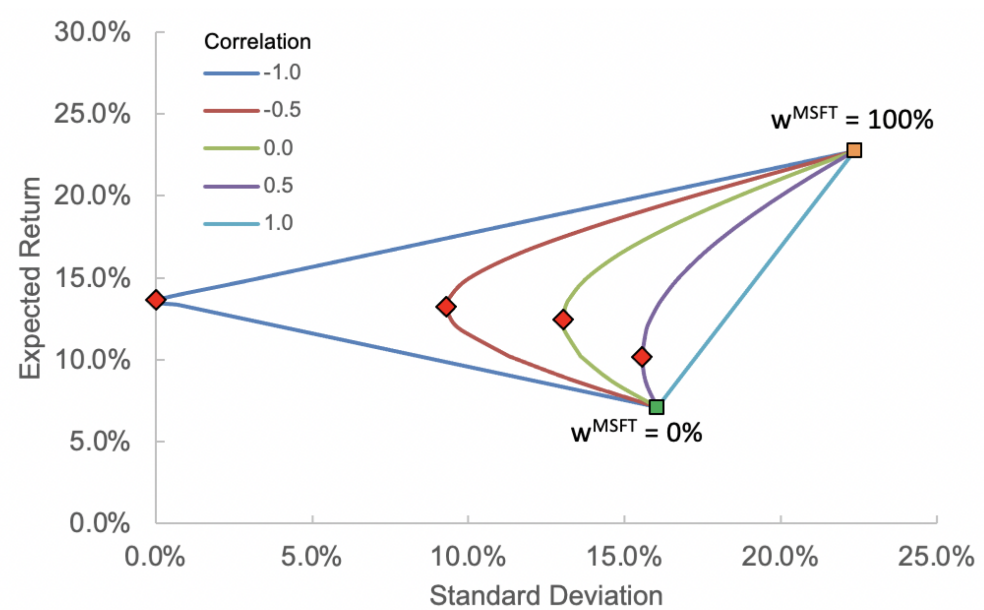

Figure 9.6: Microsoft-Walmart Mean Variance Frontiers with Different Correlations

[7]An 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.

Copyright © 2020 Michael R. Roberts

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. There is no offsetting bad performance by one stock because the two stocks always move in the same manner.

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 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]

[8] How we got 0.72 is technical and not terribly enlightening, but the accompanying Excel spreadsheet shows how to do it.

Copyright © 2020 Michael R. Roberts

9.3.4 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. To illustrate the similarities and differences, let's construct a mean variance frontier from the follow 12 stocks:

- 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 is beyond what's necessary for this book. So, let's relegate the math to the technical appendix and focus on the results and intuition.

Figure 9.7 presents the frontier and some key data points. The blue circles correspond to the average return and volatility for the 12 stocks. Each data point is 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 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.

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 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

Copyright © 2020 Michael R. Roberts

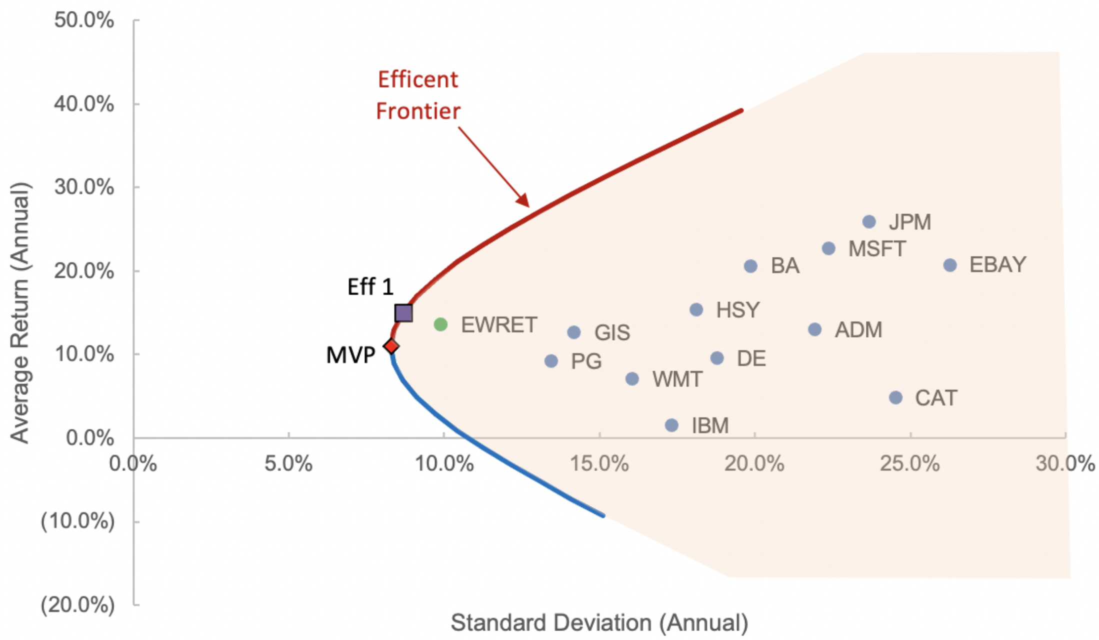

Figure 9.7: Mean Variance Frontier with Multiple Assets

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 (9.5). The key difference is that the two-asset frontier represents *all* of the portfolios that can be formed with those assets.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 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 are the efficient portfolios offering the highest expected return for a given level of risk or, equivalently, the lowest risk for a given level of expected return.

Copyright © 2020 Michael R. Roberts

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)

9.4 Adding a Risk-Free Asset

Notice in figure 9.7 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 risk as we saw in 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.[9]

Figure 9.8 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.

Portfolios consisting of the risk-free asset (T-bill) and Microsoft are represented by the

[9] Theoretically, 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.

Copyright © 2020 Michael R. Roberts

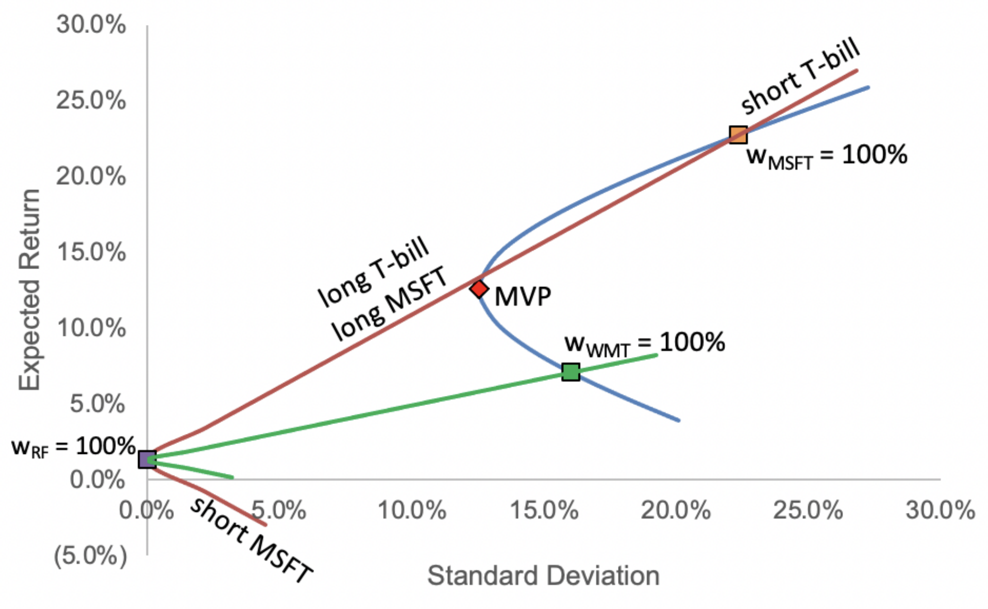

Figure 9.8: Mean Variance Frontier with Risk-Free Asset

red lines, which consists of three sections. The first is the upward sloping red line between the purple and orange square. This section represents all of the portfolios that are long the risk-free asset, Microsoft, or both (i.e., portfolio weights 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 have effectively borrowed money at the risk-free rate, combined it with any other money, and invested 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

Copyright © 2020 Michael R. Roberts

efficient portion of the frontier is different and worth discussing in some detail.

9.4.1 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 9.5). 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 9.8. 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 9.9 isolates the new efficient frontier with a risk-free asset.

This subsection refers to this efficient frontier as temporary for a reason. In fact, we can do better. Let's see how.

9.4.2 Application: Sharpe Ratio

Using equation 9.3, the expected return for our portfolio of a T-bill and Microsoft stock is

$$\mathbb{E}(r^P) = w^f \mathbb{E}(r^f) + w^{MSFT} \mathbb{E}(r^{MSFT}), \tag{9.10}$$

where $\mathbb{E}(r^P)$ is the expected return on the portfolio, and the superscript f denotes the risk-free assets. So, w^f is the portfolio weight on the risk-free asset, and $\mathbb{E}(r^f)$ is the expected

Copyright © 2020 Michael R. Roberts

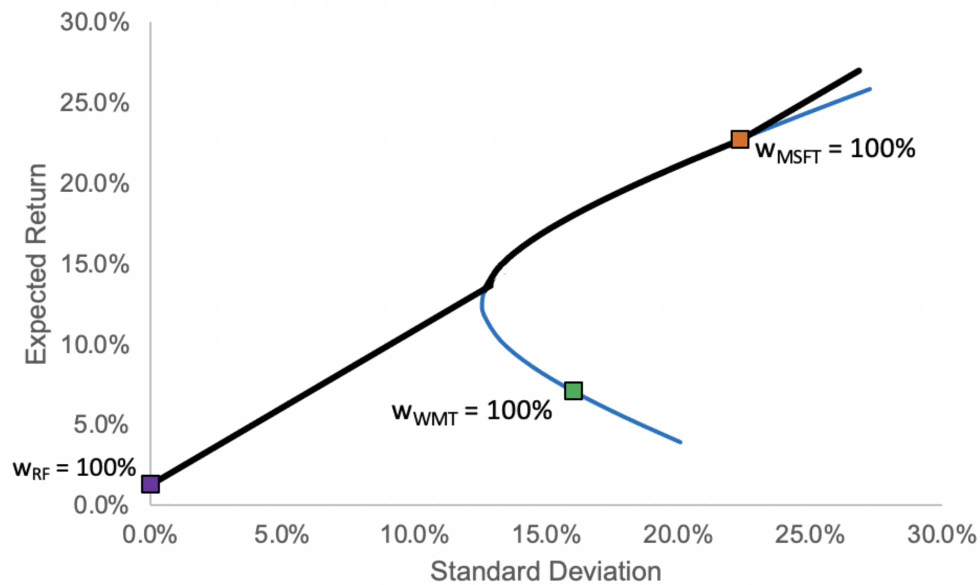

Figure 9.9: Efficient Frontier with Risk-Free Asset

return on the risk-free asset.[10] This equation is unchanged from that when we have two risky assets. We just substituted the risk-free asset's expected return. $\mathbb{E}(r^f)$ for what was Ferrari's or Walmart's.

Using equation 9.4, we can compute the volatility of our portfolio.

$$\sqrt{(w^f SD^f)^2 + (w^{MSFT} SD^{MSFT})^2 + 2w^{MSFT} w^f SD^{MSFT} SD^f Corr^{MSFT,f}} \qquad (9.11)$$

But, remember our definition of a risk-free asset - zero volatility. This assumption means that $SD^f = 0$. It also means that $CorrMSFT, f = 0$ because no matter what Microsoft's return does, the risk-free rate is always the same. That is, Microsoft's return (any asset return) and the risk-free return are uncorrelated. So, after plugging zeroes for the risk-free volatility and correlation into equation 9.11, the volatility of our T-bill-Microsoft portfolio is

$$\sqrt{w^{MSFT} SD^{MSFT})^2} = w^{MSFT} SD^{MSFT} \qquad (9.12)$$

Now, let's pay close attention to the red line in figure 9.10, 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 9.8). We can combine equations 9.10 and 9.12 to get an

[10]Technically, 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...

Copyright © 2020 Michael R. Roberts

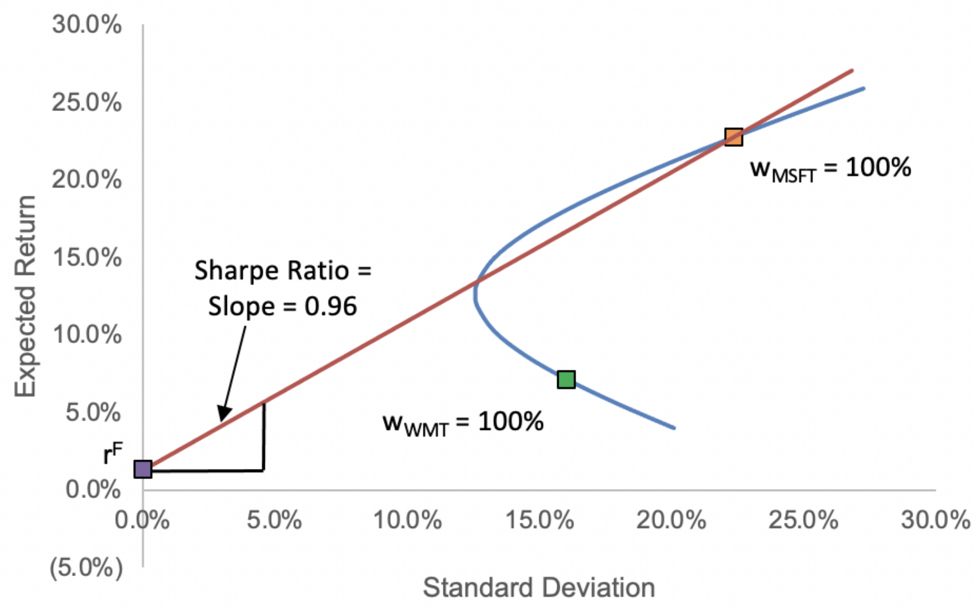

Figure 9.10: Portfolios of T-bill and Microsoft

equation for this line.

$$\underbrace{\mathbb{E}(r^P)}_{y \text{ variable}} = \underbrace{r^f}_{y\text{-intercept}} + \underbrace{\frac{\mathbb{E}(r^{MSFT}) - r^f}{SD^{MSFT}}}_{\text{slope (Sharpe Ratio)}} \times \underbrace{SD^P}_{x \text{ variable}} \tag{9.13}$$

The y variable is just the expected return on the portfolio. The y-intercept equals the risk-free rate, where the line touches the y-axis. 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 is also called a **Sharpe Ratio**, named after the Nobel price winning financial economist William Sharpe. The numerator of the Sharpe ratio measures the **risk premium** or **excess expected return** of an asset defined as the expected return minus the risk-free rate. Microsoft's risk premium is 22.73% - 1.3% = 21.43%. The denominator measures the volatility or risk of the asset - 22.36% for Microsoft. 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 9.8, 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

Copyright © 2020 Michael R. Roberts

ratio and is therefore a better investment. We can also see that portfolio of T-bills and Microsoft are strictly better than portfolios of T-bills and Walmart because they offer higher returns for the same risk.

The x variable in equation 9.13 is the portfolio volatility.

So, equation 9.13 tells the the expected return on a portfolio of a T-bill and Microsoft stock given the risk-free return, the expected return on Microsoft, the volatility of Microsoft, and the volatility of the portfolio, which we can find with equation 9.12. We can use equation 9.10 or 9.13, both will produce the same result but both provide different information.

To illustrate, assume we own a portfolio with 30% of our wealth invested in a T-bill and 70% in Microsoft. Using equations 9.10 and 9.12 tells us our portfolio expected return and volatility.

$$\begin{aligned} \mathbb{E}(r^P) &= w^f \mathbb{E}(r^f) + w^{MSFT} \mathbb{E}(r^{MSFT}) = 0.3 \times 1.3\% + 0.7 \times 22.73\% = 16.3\% \\ SD^P &= w^{MSFT} SD^{MSFT} = 0.7 \times 22.36\% = 15.65\% \end{aligned}$$

Using equation 9.13, we get the same portfolio expected return.

$$\mathbb{E}(r^P) = r^f + \frac{\mathbb{E}(r^{MSFT}) - r^f}{SD^{MSFT}} \times SD^P = 1.3\% + \frac{22.73\% - 1.3\%}{22.36\%} \times 15.65\% = 16.3\%$$

9.4.3 Application: Efficient Frontier

As said above, we like portfolios that are as far Northwest in the mean-standard deviation plane as possible - portfolios with the highest expected return and the lowest risk. Figure 9.11 shows how maximizing the Sharpe ratio helps us achieve that goal. Let's first understand the figure and how it is constructed.

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-standard deviation frontiers corresponding to different portfolios.

Take ray A for example. This ray shows the risk-reward combinations (standard deviation and expected return) for portfolios containing a T-bill and a portfolio of Microsoft and Walmart stock in which the weights are 0% and 100%, respectively. In other words, ray A corresponds to portfolios containing a T-bill and Walmart stock.

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 from the

Copyright © 2020 Michael R. Roberts

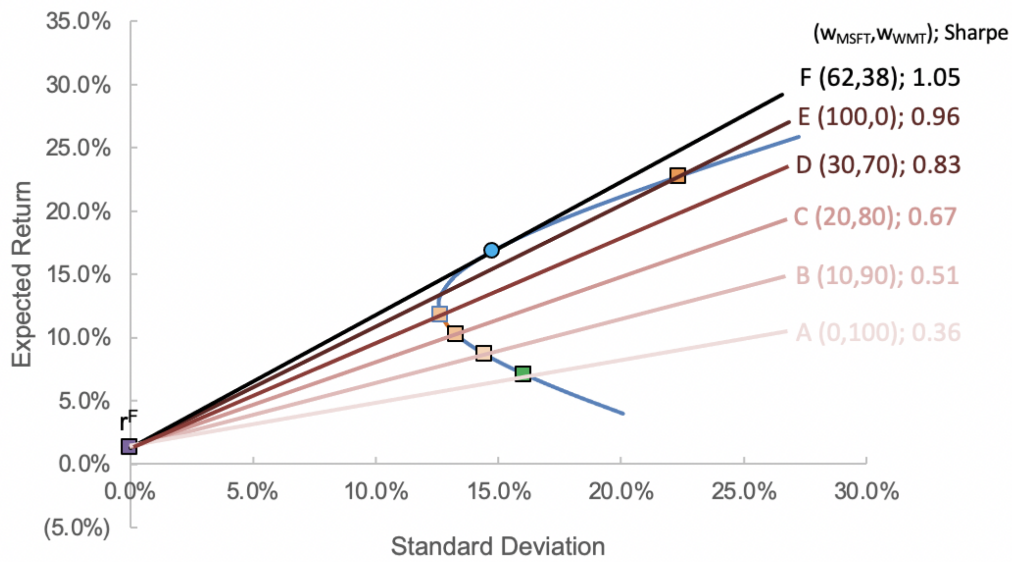

Figure 9.11: Maximizing the Sharpe Ratio

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 and is equal to $(7.1\% - 1.3\%)/16.0\% = 0.36$. Investing in Walmart stock generates 0.36 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

Copyright © 2020 Michael R. Roberts

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. This portfolio 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 computed as the expected return for the Microsoft-Walmart portfolio less the risk-free return 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 9.4. 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 which portfolio of risky assets - Microsoft and Walmart - are combined with the T-bill. The message of these different rays is that we as investors would like to be on the ray with the steepest slope, i.e., the highest Sharpe ratio. Why? Pick a point on any other ray. Notice that for the same risk, we can find a portfolio with a higher expected return by moving vertically up to a higher ray. Alternatively, for the same expected return, we can find a portfolio with lower risk.

The steepest ray is F. This ray is **tangent** to the blue curve meaning it just barely touches the curve at a point. The point at which it touches the curve, indicated by the turquoise circle, is called the **tangency portfolio**. In this example, the tangency portfolio has weights of 62.1% and 37.9% on Microsoft and Walmart (details in the technical appendix).

There are no rays more steeply sloped than F because there are no investment opportunities outside of the blue curve. Notice that we can't "improve" on any portfolios found on ray F. If we want a higher return than a portfolio on ray F offers, we have to accept more risk (move up and to the right along the ray). If we want less risk we have to accept a lower return (move down and to the left along the ray). The punchline here is that ray F is our new and final efficient frontier when we are only considering three assets - the risk-free security, Walmart stock, and Microsoft stock.

Copyright © 2020 Michael R. Roberts

Importantly, this conclusion holds true when we consider more than three assets. Figure 9.12 presents the efficient frontier for the 12 stocks discussed above in section 9.3.4. Notice the similarity to figure 9.12. 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 on 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.

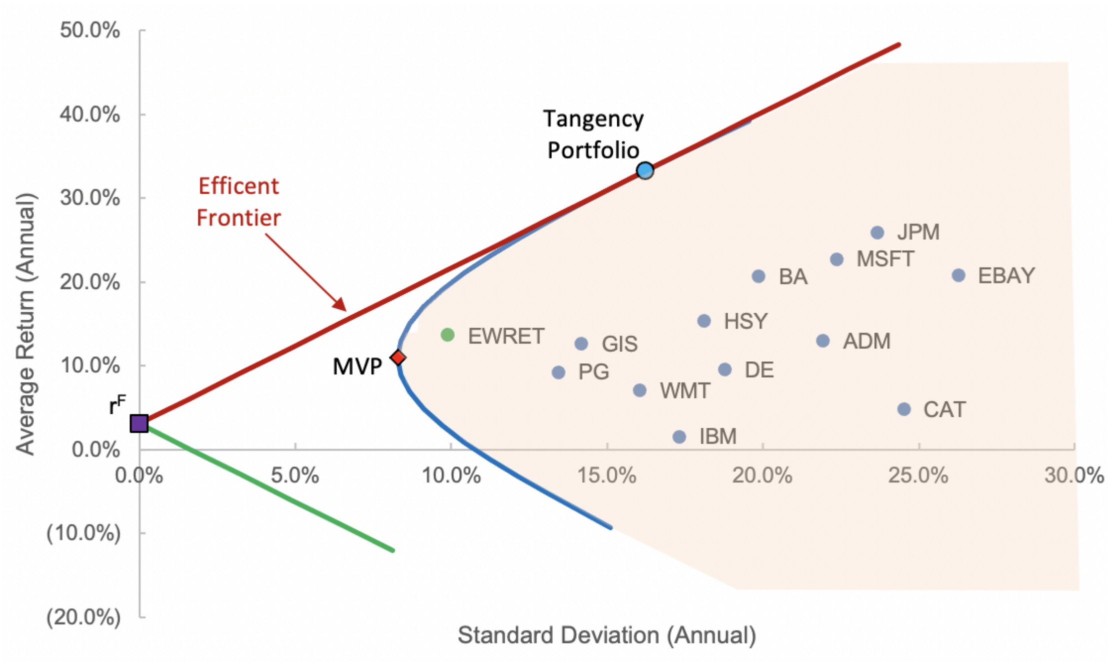

Figure 9.12: Efficient Frontier for 12 Stocks and Risk-free Asset

9.4.4 Application: 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, the interest rate on our mortgage is higher than the rate on our savings account or on a Treasury security.

Copyright © 2020 Michael R. Roberts

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

Each ray in figure 9.11 assumes the that these rates are the same. This assumption can be seen in equation 9.13, 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 9.13 shows what happens to these rays when there are different borrowing and lending rates.

The blue curve is our usual risk-asset (Microsoft and Walmart) frontier containing the tangency portfolio as indicated by the turquoise circle. The purple square on the vertical axis is our risk-free return of 1.3%. However, now we can only lend (i.e., invest) at this rate, hence the notation r^{LEND}. If we want to borrow money, it will cost us 8% as indicated by the green square on the vertical axis and 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 therefore represents only hypothetical portfolios.

The portfolios on the dashed red line are similarly hypothetical. We can't lend at 8%,

Copyright © 2020 Michael R. Roberts

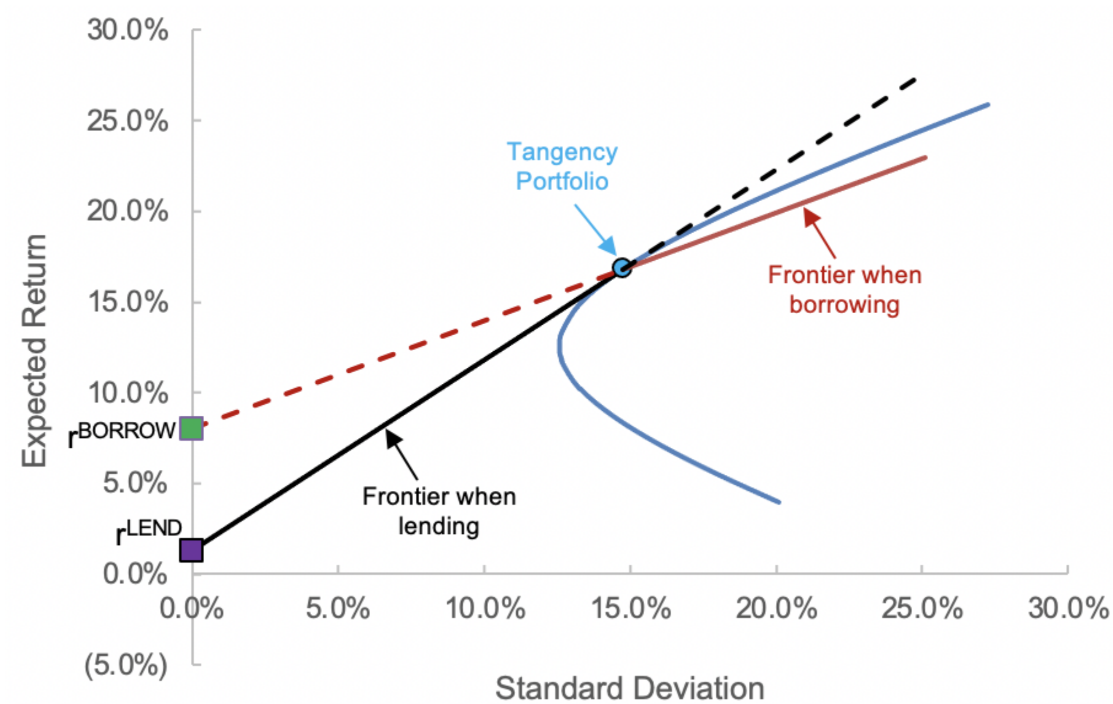

Figure 9.13: Efficient Frontier for 12 Stocks and Risk-free Asset

i.e., earn a risk-free return of 8%. (Oh, but I wish we could today in late 2021.). 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 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.

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

$$\mathbb{E}(r^P) = r^{LEND} + \frac{(r^T) - r^{LEND}}{SD^T} SD^P.$$

When we are short the risk-free asset ($w^F < 0$), we are borrowing money and the equation for the solid red line is

$$\mathbb{E}(r^P) = r^{BORROW} + \frac{(r^T) - r^{BORROW}}{SD^T} SD^P.$$

These two equations express the expected return on the portfolio containing the risk-free asset and the tangency portfolio. They are similar to equation (9.13) with two differences.

Copyright © 2020 Michael R. Roberts

First, we are now considering portfolios of the risk-free asset and the tangency portfolio instead of Microsoft. Second, depending on whether we are long (lend) or short (borrow) the risk-free asset, we have a different risk-free return.

The main message of this analysis is that shorting the risk-free asset to increase returns, even if we're willing to stomach the greater volatility, is no longer efficient because shorting comes at a higher cost ($r^{BORROW} > r^{LEND}$).

9.5 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 to avoid 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 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.

Copyright © 2020 Michael R. Roberts

These takeaways can also be summarized with two pictures shown in figures 9.14 and 9.15. In both figures, the horizontal axis measures the standard deviation of returns, the vertical axis measures average or expected return.

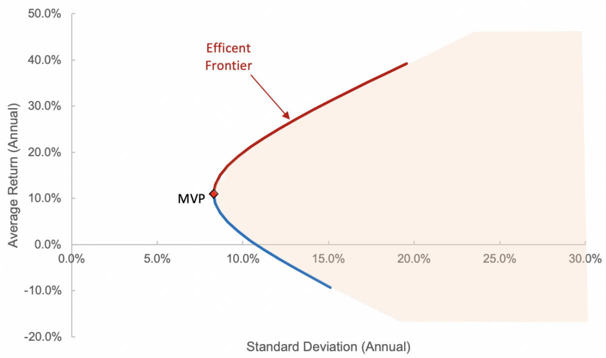

Figure 9.14: Mean Variance and Efficient Frontier of Risky Assets

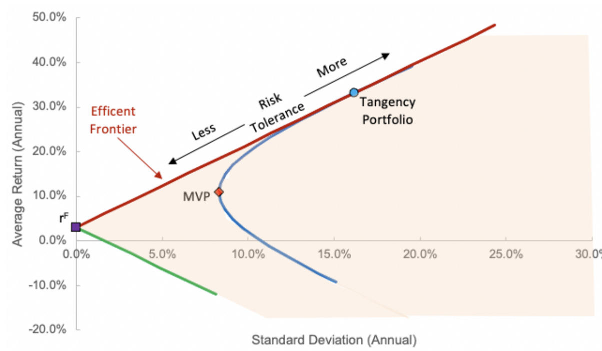

Figure 9.15: Mean Variance and Efficient Frontier of Risky and Risk-free Assets

In figure 9.14, 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 9.15 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 (thousands) assets? In other words, to exploit what we've learned in

Copyright © 2020 Michael R. Roberts

this chapter, we need to know the tangency portfolio for *all* assets in which we can invest. That's next.

9.6 Technical Appendix

More details on the math and some results from this chapter are presented here.

9.6.1 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 9.9. 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

$$\frac{\mathbb{E}(r^P) - r^F}{SD^P} = \frac{w^A \mathbb{E}(r^A) + (1 - w^A)\mathbb{E}(r^B) - r^F}{\sqrt{(w^A)^2(Var^A)^2 + (1 - w^A)^2(Var^B)^2 + 2w^A(1 - w^a)Cov^{A,B}}}$$

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

Copyright © 2020 Michael R. Roberts

for w^A. Computing the derivative is tedious but just repeated application of the chain rule. The result of this process is as follows.

$$\begin{aligned} w^A &= \frac{(\mathbb{E}(r^A) - r^F)Var^B - (\mathbb{E}(r^B) - r^F)Cov^{A,B}}{(\mathbb{E}(r^A) - r^F)Var^B + (\mathbb{E}(r^B) - r^F)Var^A + (\mathbb{E}(r^A) - r^F + \mathbb{E}(r^B) - r^F)Cov^{A,B}} \\ w^B &= 1 - w^A \end{aligned}$$

9.6.2 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

$$w = (w^1, ..., w^N), \text{ and}$$
$$\mu = \left(\mathbb{E}(r^1), ..., \mathbb{E}(r^N)\right].$$

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

$$w'\mu = \mathbb{E}(r^P) = w^1\mathbb{E}(r^1) + ... + w^N\mathbb{E}(r^N) \tag{9.14}$$

The variance of our N-asset portfolio is similar to that of our 2-asset portfolio (equation 9.5), just with more variance and covariance terms.

$$\begin{aligned} w'\Sigma w = Var^P &= w_1^2\sigma_1^2 + ... + w_N^2\sigma_N^2 \\ &+2w_1w_2\sigma_{1,2} + ... + 2w_1w_N\sigma_{1,N} \\ &+2w_2w_3\sigma_{2,3} + ... + 2w_2w_N\sigma_{2,N} \\ &+... \\ &+2w_{N-1}w_N\sigma_{N-1,N} \end{aligned} \tag{9.15}$$

Copyright © 2020 Michael R. Roberts

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.

$$\begin{aligned} \min_{w} \quad & \frac{1}{2} w' \Sigma w \\ \text{s.t.} \quad & \mu' w \geq m \\ & e' w = 1 \end{aligned} \tag{9.16}$$

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., $e = (1, ..., 1)$).

$$e' w = w^1 + w^2 + ... + w^N$$

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.

$$w^{Efficient} = (1 - a) w^1 + a w^2$$

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.

$$\begin{aligned} \min_{w} \quad & w' \Sigma^{-1} w \\ \text{s.t.} \quad & e' w = 1 \end{aligned}$$

Copyright © 2020 Michael R. Roberts

The solution to this program is

$$w_{MVP} = \left(e'\Sigma^{-1}e\right)^{-1}\Sigma^{-1}e.$$

Plugging this solution into the equations 9.14 and 9.15, we obtain the expected return and variance of the minimum variance portfolio, respectively.

$$E(r^{MVP}) = \frac{\mu\Sigma^{-1}e}{e'\Sigma^{-1}e} \quad (9.17)$$

$$Var^{MVP} = \sigma_{mvp} = e'\Sigma e \quad (9.18)$$

Tangency Portfolio

The tangency portfolio is the portfolio of risky assets with the maximum Sharpe ratio and therefore solves the following program.

$$\max_{w} \frac{\mu' w - r^F}{w'\Sigma w}$$
$$\text{s.t.} \quad e'w = 1$$

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 identif

$$w^{Efficient} = (1 - a)w^{MVP} + aw^{Tang}$$

For any target return, m, a equals

$$\frac{m\left(\mu'\Sigma^{-1}e\right)\left(e'\Sigma^{-1}e\right) - \left(\mu'\Sigma^{-1}e\right)^2}{\left(\mu'\Sigma^{-1}\mu\right)\left(e'\Sigma^{-1}e\right) - \left(\mu'\Sigma^{-1}e\right)^2} \quad (9.19)$$

Copyright © 2020 Michael R. Roberts

9.7 Problems

9.1 TBD

Copyright © 2020 Michael R. Roberts

Copyright © 2020 Michael R. Roberts

Part III

Decisions Finance Specialists Make

Copyright © 2020 Michael R. Roberts

Chapter 10

Estimating the Cost of Capital - The CAPM and Multifactor Models

---

Fundamental value relation

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

This chapter

- presents a model - the capital asset pricing model - to estimate expected returns - the discount rate r in our fundamental valuation relation,
- decomposes risk into two components - that which can and that which cannot be eliminated by investors diversifying their investments,
- argues that under certain circumstances people should simply invest their money in a broad-based, low-cost index fund and treasury securities,
- derives an important measure of investment performance,
- introduces factor models more generally and their approach to estimating expected returns, and
- applies our fundamental value relation to answer several questions including:
  - What is Apple Inc.'s cost of capital?

Copyright © 2020 Michael R. Roberts

– How does the risk of Apple, or other assets, compare to the of the broader stock market? Treasury securities?

– Has investing in Apple produced returns consistent with its risk?

Estimating the cost of capital, r in our fundamental value relation, is nothing new. For example, we've estimate expected stock returns using:

- arithmetic averages of historical returns,
- the dividend discount model, and
- the total payout model

in chapter 8. Likewise, we estimated expected returns for defaultable bonds in the technical appendix of chapter 7. (The expected return for a non-defaultable bond is just the yield on the bond.)

This chapter provides an alternative approach based on economic models. We'll focus on one in particular, the Capital Asset Pricing Model or CAPM, that is widely used in practice. We'll mention some alternatives at the end. The **Capital Asset Pricing model** or **CAPM** is a theory of how investors behave and the implications of that behavior for expected returns. I'm not going to do this Nobel Prize winning theory justice here.[1] This was one of those hard decisions I mentioned in the Preface. Instead, I'll focus on the key insights from the theory of relevance to decision makers.

A warning: The CAPM, like all economic models, is wrong in that it vastly simplifies reality and has questionable empirical validity. 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 Everyone Should Invest in the Market Portfolio

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. To get this result, the CAPM relies on the following assumptions.

[1] William Sharpe (as in Sharpe ratio) won the 1990 Nobel Prize for formulating the Capital Asset Pricing Model.

Copyright © 2020 Michael R. Roberts

1. Investors only care about expected returns and volatility.

2. Investors share the same beliefs about how asset returns will evolve going forward.

3. Investors can borrow and lend at the same risk-free rate.

Assumption 1 is inaccurate but not crazy. Most investors worry about the return they'll earn and the volatility of that return - how much it varies from period to period. However, some (many) investors worry about other features of returns such as the likelihood of a large loss. Assumption 2 is crazy. Most people have different views about how different stocks, bonds, etc. will perform in the future. Assumption 3 is inaccurate and partially crazy. Most people can buy a Treasury security or some other relatively safe and reliable asset, but there is no such thing as a literal risk-free asset (i.e., one with no price volatility) and good luck trying to borrow at the same rate as the U.S. federal government. Grain of salt...

These assumptions imply 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.

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.

This generates a key insight of the CAPM: Investors should follow a simple investment strategy. Put all of our savings into a combination of a risk-free asset, like a Treasury security, and the market portfolio, like Vanguard's Total Stock Market Portfolio. 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

Copyright © 2020 Michael R. Roberts

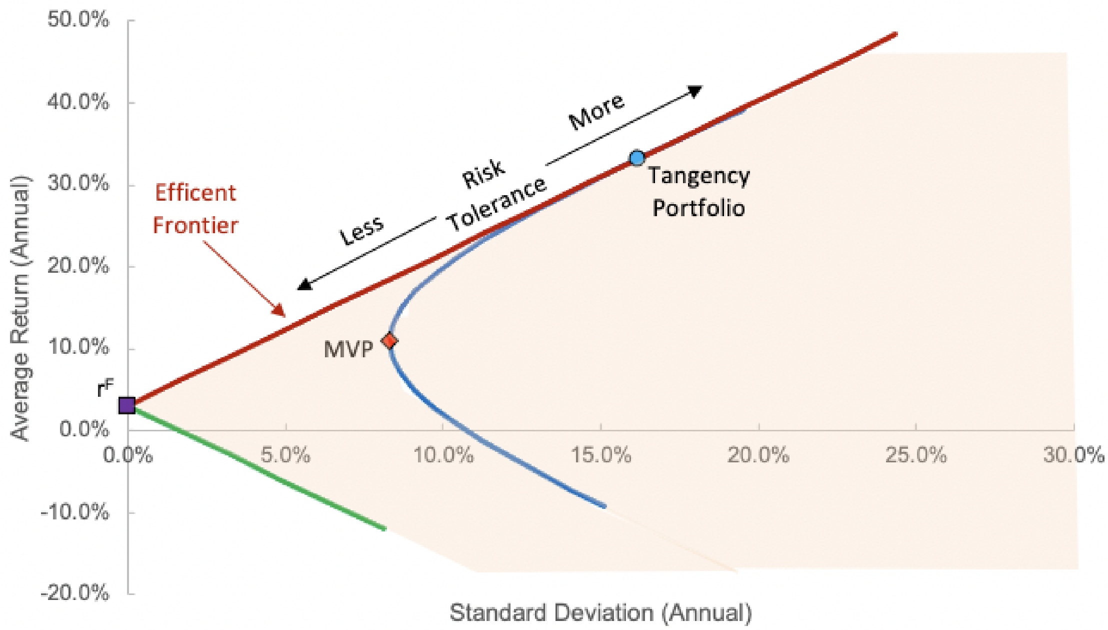

Figure 10.1: Mean Variance and Efficient Frontier of Risky and Risk-free Assets

CAPM world. And, the only thing that might change over time is how much we invest in the risk-free asset and the market to reflect changes in our risk tolerance.

Now, as we mentioned, there is no such thing as a risk-free asset in a CAPM sense. A Treasury security is just a proxy. Likewise, stock market mutual funds do not represent "the market" in a CAPM sense, which includes stocks, bonds, real estate, and all other assets. We could invest in additional assets, such as bonds, real estate, venture capital, etc., in addition to a market index to further diversify our holdings. Ultimately anything we do will be a proxy for the CAPM's notion of the market because there are so many assets in the economy that are not tradeable (e.g., privately held companies).

Despite these limitations and some unrealistic assumptions, the CAPM is simple and practical, which is why it is so popular.

10.2 Estimating Expected Returns

Previously we estimated expected returns, $\mathbb{E}(r)$, with a simple average,

$$\bar{r} = \frac{1}{T}(r_1 + ... + r_T)$$

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 returns, $\mathbb{E}(r)$. It was our best guess of what an asset,

Copyright © 2020 Michael R. Roberts

like a bond or stock, would earn in the future.

The CAPM provides an alternative approach for estimating expected returns.

$$\mathbb{E}(r) = r^F + \beta\left(\mathbb{E}(r^M) - r^F\right) \tag{10.1}$$

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**, $\beta\left(\mathbb{E}(r^M) - r^F\right)$. This term represents the additional compensation an investor will earn when taking risk; hence, the name risk premium.

Note that we wrote r^F instead of $\mathbb{E}(r^F)$. We could have written either because the risk-free return, at least in theory, never varies; it's standard deviation equals zero. This means r^F is just one number (e.g., 5%). So, our best guess of what the risk-free rate next year is just the number it will always be; $\mathbb{E}(r^F) = r^F$. More generally, the expectation of a constant is that constant.

The risk-free return is in contrast to the expected return on a risky asset, like a stock, whose realized return is changing from day to day, month to month, and year-to-year (i.e., standard deviation is greater than zero). As we've previously discussed, our best guess of the stock's return next year, $\mathbb{E}(r)$, may be very different from the stock's realized return next year, r.

Risky assets earn a premium over the risk-free return, and this risk-premium differs across assets. Specifically, 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.

2. **Market risk premium**, $\left(\mathbb{E}(r^M) - r^F\right)$. This term represents the compensation investors receive for investing in the market. It is also referred to as the **market risk factor** and **excess market 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

Copyright © 2020 Michael R. Roberts

larger an asset's beta, the greater the asset's exposure to market fluctuations, the greater the risk, and the greater the return. So, beta is a new measure of risk - an alternative to standard deviation (a.k.a., volatility).

The questions now are: What exactly is β, and how do we use equation 10.1 in practice? In other words, how do we put numbers to r^F, β, and $\mathbb{E}(r^M)$ to get a number for $\mathbb{E}(r)$, the expected return we're interested in estimating?

10.3 Application: 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.

10.3.1 Step 1: Estimate the Risk-Free Rate

There is no true 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. But, which one? A T-bill? T-note? T-bond? And, what maturity? 30-day? 5-year? 30-year?

The choice of 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 they 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 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 the 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.

10.3.2 Step 2: Estimate the Market Risk Premium

The market risk premium, $\mathbb{E}(r^M) - r^F$, 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

Copyright © 2020 Michael R. Roberts

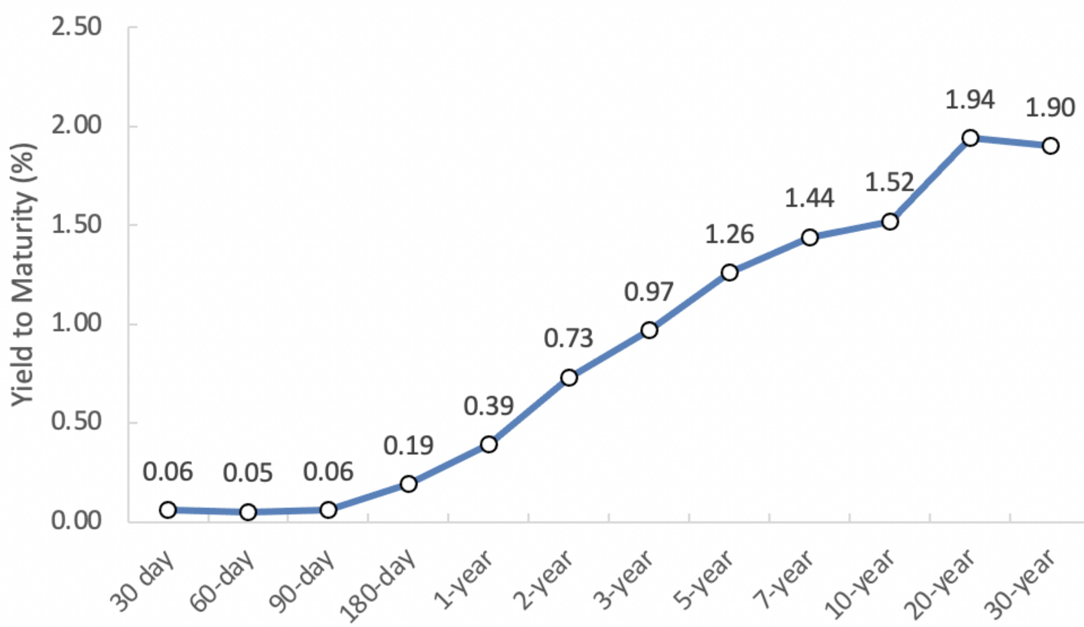

Figure 10.2: March 19, 2021 Treasury Yield Cure (Source: U.S. Treasury)

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 take an average of historical returns to estimate the expected return on the market, we can take an average of historical returns to estimate the expected return on the Treasury.

	Returns (%)	
Investment	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.[2] 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

[2]The U.S. stock market is a value weighted portfolio of all stocks on the NYSE, AMEX, and NASDAQ.

Copyright © 2020 Michael R. Roberts

1970 to 2021, those estimates vary from 7.87% to 5.09%. We'll use 6% for the purpose of this application.

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.

$$\beta^A = \frac{Cov^{A,M}}{Var^M} \tag{10.2}$$

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 returns 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.[3] From January 2017 to December 2021, Apple's market beta was.

$$\beta = \frac{0.00232}{0.00206} = 1.13$$

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, Apple increases by 1.13%.

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.

$$\mathbb{E}(r^{Apple}) = r^F + \beta^i \left(\mathbb{E}(r^M) - r^F\right) = 0.019 + 1.13 \times 0.06 = 0.0868$$

[3] Another common choice is the most recent two years of weekly data.

Copyright © 2020 Michael R. Roberts

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% is awfully low. In fact, from 2017 to 2021, Apple's average *monthly* return was 3.52%. It's average annual return was 45.32%. So, it's natural to ask: 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.

$$\mathbb{E}(r^i) = r^F + 1 \times \left(\mathbb{E}(r^M) - r^F\right) = \mathbb{E}(r^M)$$

It can also be seen in equation 10.2. The beta for the market portfolio, β^M is

$$\beta^M = \frac{Cov^{M,M}}{Var^M} = \frac{Var^M}{Var^M} = 1. \tag{10.3}$$

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.

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 *in*sensitive to market fluctuations and, as such, offer lower returns than the market, on average.

Copyright © 2020 Michael R. Roberts

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.

$$\mathbb{E}(r^i) = r^F + 0 \times \left(\mathbb{E}(r^M) - r^F\right) = r^F$$

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 it's 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.[4] So, 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. If this sounds odd, that's ok. Let's explore it.

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 your 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 reduced by diversifying your 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.

[4] Remember from chapter 9 that two returns that are uncorrelated (i.e., zero correlation) have zero covariance.

Copyright © 2020 Michael R. Roberts

Systematic risk is captured by the risk premium term in equation 10.1, $\beta^i\left(\mathbb{E}(r)-r^F\right)$. Idiosyncratic risk is a little more subtle.

Let's remove the expectation operators - the $\mathbb{E}$s - from equation 10.1 and add an "error" term, e, that is uncorrelated with the market return (r^M) and whose expectation is zero $(\mathbb{E}(e)=0)$.

$$r^i = r^f + \beta^i\left(r^M - r^F\right) + e \tag{10.4}$$

This equation says that asset i's return, r^i, is equal to the risk-free rate plus the risk premium plus an error. The terms r^i, r^M, and e are all random variables that can take on many different values with different probabilities. If we take expectations of both sides of this equation, we recover equation 10.1 because our best guess for the 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.

$$Var^i = \underbrace{(\beta^i)^2 Var^M}_{\text{Market risk}} + \underbrace{Var^e}_{\text{Idiosyncratic risk}} \tag{10.5}$$

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.

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.

$$Var^e = Var^i - (\beta^i)^2 Var^M = 0.0072 - (1.13)^2 \times 0.0021 = 0.0046$$

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.

Copyright © 2020 Michael R. Roberts

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. Because of this, the risk that comes from variation in this error term - the idiosyncratic risk - is “not priced” and is referred to as **unpriced** risk. It is risk that does *not* affect expected returns. Investors are not compensated for firm-specific risk because, according to the CAPM, they can easily remove this risk by holding diversified portfolios.

In contrast, the expected market return *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, they cannot get rid of the risk that affects *all* stocks.

10.4.4 Application: Beta Variation

How much do betas vary over time and across assets? Figure 10.3 shows Apple and Advanced Micro Devices (ticker=AMD) 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.[5]

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, 2.1, means a very big difference in the equity cost of capital for AMD.

Assume interest rates and market risk premium are constant at 1.9% and 6%. 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.

[5] In 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.

Copyright © 2020 Michael R. Roberts

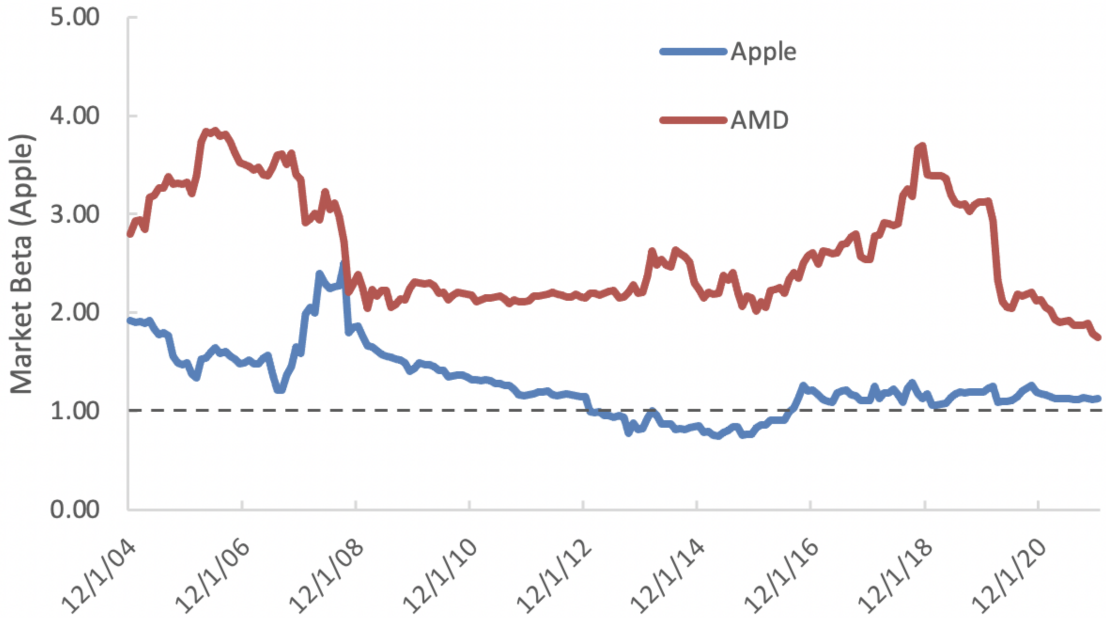

Figure 10.3: Time Variation in Market Betas

10.5 Bringing it all Together

Figure 10.4 brings most everything we've learned in the previous and current chapters together. The figure is comprised of two plots and a lot of dashed lines connecting them.

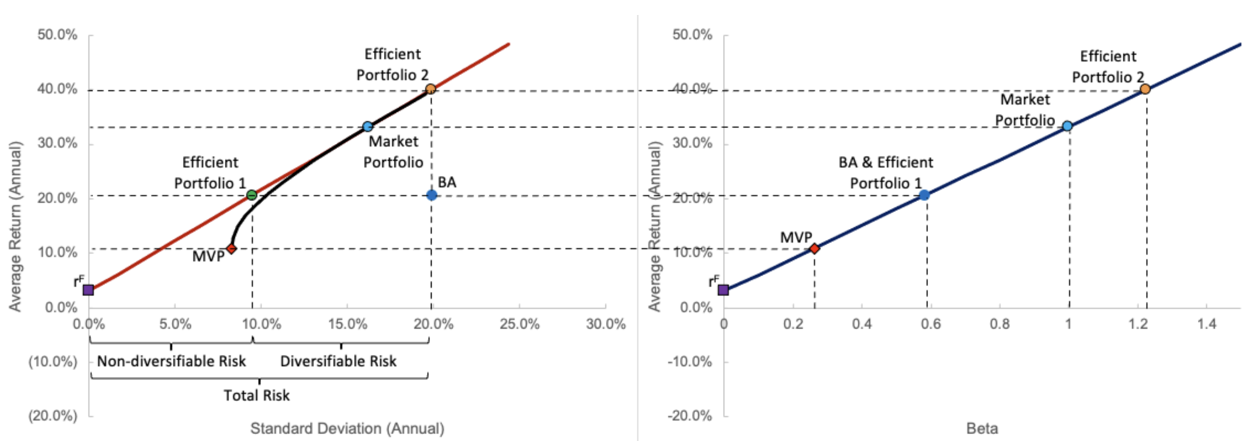

Figure 10.4: Capital Market Line (CML) and Security Market Line (SML)

Copyright © 2020 Michael R. Roberts

10.5.1 Capital Market Line and Efficient Portfolios

The left plot in figure 10.4 presents the mean-variance frontier of 12 risky assets (black curve) and the efficient frontier (red ray) that we analyzed in chapter 9.[6] The horizontal axis measures standard deviation; the vertical axis measures expected returns. We estimate both statistics using the standard deviation and average of historical returns. The figure also highlights several specific portfolios, where the term portfolio is used to 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.
- Efficient portfolio 2 - orange circle on the red ray.

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 of 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.

[6]The 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).

Copyright © 2020 Michael R. Roberts

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%. The standard deviation of any portfolio on the efficient frontier equals the weight on the market times the standard deviation of the market. Remember that the efficient frontier contains portfolios of the risk-free asset, whose return doesn't vary, and the market. To find the weight on the market, we can set the volatility of the efficient portfolio equal to Boeing's volatility and solve for the weight like so.

$$SD^{BA} = 0.199 = w^M SD^M = w^M 0.162 \implies w^M = \frac{0.199}{0.162} = 1.22$$

If we are willing to accept a volatility of 19.9% year, we would be much better off investing 122% of our wealth in the market and -22% in the risk-free asset (i.e., borrow money) than investing in Boeing stock alone. In fact, the expected return to efficient portfolio 2 is

$$w^F r^F + w^M \mathbb{E} r^M = -0.22 \times 0.031 + 1.22 \times 0.3325 = 0.3988,$$

or 39.88%, significantly higher than the 20.6% offered by Boeing.

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.

Copyright © 2020 Michael R. Roberts

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 peice 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%
- Systematic risk = 9.4% (We computed the volatility of Efficient Portfolio 1 just above.)
- Idiosyncratic risk = 11.2% (20.6% - 9.4%)

Investing in Boeing *alone* is inefficient in this setting beacause 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.

$$\mathbb{E}(r^P) = r^F + \underbrace{\frac{\mathbb{E}(r^M) - r^F}{SD^M}}_{\text{Market Sharpe Ratio}} SD^P \tag{10.6}$$

Copyright © 2020 Michael R. Roberts

The left side of the equation is the expected return on an efficient portfolio. The intercept is the risk-free return, r^F. The x-variable is the standard deviation of an efficient portfolio, SD^P. The slope is the Sharpe ratio of the market portfolio, whose value 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 10.4 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 horitzontal 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. It's worth repeating.

$$\mathbb{E}(r^i) = r^F + \beta^i \left(\mathbb{E}(r^M) - r^F\right)$$

The security market line says that *every* 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, $\left(\mathbb{E}(r^M) - r^F\right)$.

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.

The beta of a portfolio equals the weighted sum of the individual betas where the weights are the portfolio weights. In other words, the beta of a portfolio with N assets is

$$\beta^P = w^1\beta^1 + w^2\beta^2 + ... + w^N\beta^N, \tag{10.7}$$

Copyright © 2020 Michael R. Roberts

where $w^1, w^2, ..., w^N$ are the portfolio weights and $\beta^1, \beta^2, ..., \beta^N$ are the individual asset betas.

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 10.7, 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 in this example.

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.

10.6 Application: Jensen's Alpha

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. There is another way to estimate beta that results in a similar estimate. It begins with equation 10.4, which is repeated here and is referred to as a **population regression function**.

$$r^i = r^f + \beta^i \left(r^M - r^F \right) + e$$

The random variable, asset i's return r^i, is a linear function of other random variables, the market return r^M and the error term e, plus an intercept equal to the risk-free rate.

Let's subtract the risk-free rate from both sides of the equation and add an intercept term, α.

$$\underbrace{\left(r^i - r^f \right)}_{y} = \alpha + \beta^i \underbrace{\left(r^M - r^F \right)}_{x} + e$$

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, $(r^M - r^F)$. The intercept, α, should equal zero according to the CAPM.

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

Copyright © 2020 Michael R. Roberts

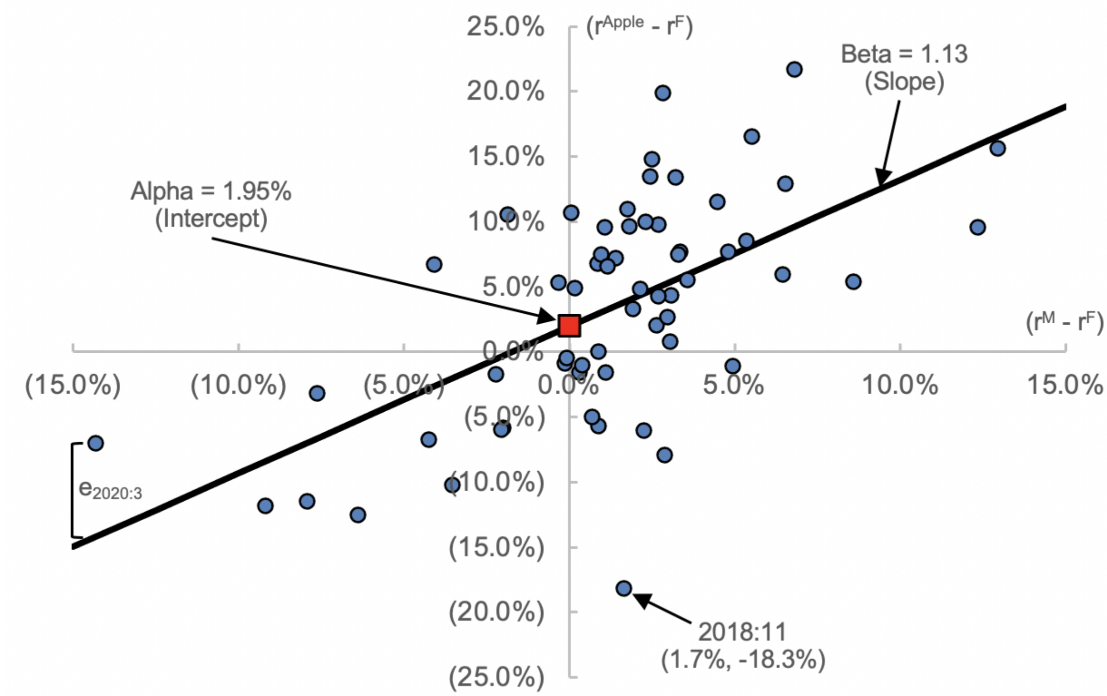

Figure 10.5: Apple and Market Excess Returns and Estimated Regression Function

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 10.5. 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.

The slope of the regression line, 1.13, is Apple's estimated market beta, β. The y-intercept indicated by the red square is the estimated alpha, which equals 1.95% *per month*. 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 asset management. Jensen's alpha represents **excess** or **abnormal** return. It is the component of returns *not* explained by risk or, more precisely in this case, market risk.

Copyright © 2020 Michael R. Roberts

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 an *additional* 1.95% per month (23.43% per annum). Why?

Two possible reasons why Apple did so well. 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 ommitted risk factors. Second, Apple is mispriced. In other words, the market has been too optimistic about Apple's prospects (overestimated future cash flows) or underestimated the future risk (i.e., discount rate). Because it's returns are higher than what the CAPM predict, Apple may be **overvalued** today.

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 overvalued. We could short the stock in the expectation that the price will fall. But, when will the price fall? In other words, the market has to agree with us at some point so that the price on Apple actually does fall. As the famous economist, John Maynard Keynes, once said: "the market can stay irrational longer than you can stay solvent."

10.7 Other Asset Pricing 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. However, these alternative models have not gotten much use in estimating the cost of capital. Why? They often produce strange estimates (e.g. 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.[7]

$$\mathbb{E}(r^i) = r^F + \beta^M \underbrace{\left(\mathbb{E}(r^M) - r^F\right)}_{\text{Market factor}} + \beta^{SMB} \underbrace{\left(\mathbb{E}(r^{Small}) - \mathbb{E}(r^{Big})\right)}_{\text{Size (SMB) factor}} + \beta^{HML} \underbrace{\left(\mathbb{E}(r^{High}) - \mathbb{E}(r^{Low})\right)}_{\text{Value-growth (HML) factor}} \quad (10.8)$$

This model adds to additional risk factors to the market risk factor from the CAPM (equation 10.1).

[7]This 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.

Copyright © 2020 Michael R. Roberts

1. **Size factor**. This factor was motivated by the observation that smaller 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. 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 - $\left(\mathbb{E}(r^{Small}) - \mathbb{E}(r^{Big}\right)$ - based on data from 1926 to 2021 is 2.37%.

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 - after accounting for differences in market betas *and* firm size across the two groups. Book-to-market-equity is simply 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 - $\left(\mathbb{E}(r^{High}) - \mathbb{E}(r^{Low}\right)$ - based on data from 1926 to 2021 is 4.04%.

The intuition behind the Fama-French 3-factor model is simple. Market risk exposure does a poor job of explaining variation in expected returns because there are many assets with significant alpha - positive and negative. So, there must be other risks in the economy to which firms 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 10.8 eliminates alpha for many assets.

10.7.1 Application: 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.

Copyright © 2020 Michael R. Roberts

Estimate Risk Premia

This step requires us to estimate not just the market risk premium, $\left(\mathbb{E}(r^M) - r^F\right)$, 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 10.8, we need to estimate a multivariate regression using historical data. Specifically, we estimate the following regression using 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 2: 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 - 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 return 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.

Copyright © 2020 Michael R. Roberts

Putting it all together

With our beta and risk premia estimates, we can compute the expected return to Apple using equation 10.8.

$$E(r^{Apple}) = 0.019 + 1.33 \times 0.06 - 0.36 \times 0.0237 - 0.76 \times 0.0404 = 0.0595$$

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, and even before then, there have been by some accounts hundreds of other factors thought to be responsible for expected returns. Campbell Harvey and Yan Liu document over 400 factors found by academics. (See Harvey and Liu, 2020, "A Census of the Factor Zoo," Working Paper, Duke University.) Why so many? Lots of reasons, though data-mining, p-hacking, and non-results are popular explanations for many of them.

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.

- 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 Vanguards total stock market fund or S&P 500 fund.

Copyright © 2020 Michael R. Roberts

- 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.

$$\mathbb{E}(r^P) = r^F + \underbrace{\frac{\mathbb{E}(r^M) - r^F}{SD^M}}_{\text{Market Sharpe ratio}} \times SD^P$$

- 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 - accroding to the CAPM.

- The security market line relates beta, our measure of market risk exposure, to expected returns for all assets.

$$\mathbb{E}(r^i) = r^F + \beta^i \left(\mathbb{E}(r^M) - r^F\right)$$

- 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 Problems

10.1 TBD

Copyright © 2020 Michael R. Roberts

Chapter 11

Optimal Financial Policy

Copyright © 2020 Michael R. Roberts

Chapter 12

Corporate Valuation

Copyright © 2020 Michael R. Roberts

Appendix A

Financial Accounting

This chapter is a detour into accounting *for those that need it*. Financial information is stored and reported in financial statements (e.g., income statement, balance sheet, cash flow statement). The concepts discussed here are useful on their own and for deepening our understanding of some concepts in the book (e.g., capital budgeting, stock valuation, corporate valuation). We're not going to learn how to construct financial statements. That's what accountants do. We're going to learn how to use the statements.

We'll use Microsoft as our illustrative vehicle again. Microsoft operates on a June fiscal year end meaning the statements cover the period July 1, 2020 to June 30, 2021. A reminder: Don't misinterpret the discussion here as only applying to large, publicly-traded, tech companies. Financial statements for all companies are similar and their analysis proceeds along similar lines. What differs are the numbers and the story they tell.

A.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**. This statement tells us how the company made and spent money during a period, typically a quarter or a year. It can also be used for divisions, even projects, within a company to identify the sales and expenses of that division or project. 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 names or **line items** in their income statements, even if the words mean the same thing. While there is a consistent theme to all income statements, this inconsistency can make cross-company comparisons difficult.

Copyright © 2020 Michael R. Roberts

($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)

In response, financial analysts developed a common income statement format that consists of several key metrics and that can be used 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, loosely speaking, represents money coming into the firm. We'll see that sales and cash inflows do not always coincide.

2. Subtract an expense (negative expenses correspond to income), which, loosely speaking, represents money going out of the firm. We'll see that expenses and cash outflows also do not always coincide.

3. Compute a measure of earnings (or income).

Copyright © 2020 Michael R. Roberts

($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 & adminstrative 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)

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.

To draw an analogy, 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 quite similar. A sometimes useful exercise to strengthen intuition and personal budgeting is to draw parallels between the corporate P&L on which we'll focus and your personal P&L.

1.1.1 Sales

Every 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

Copyright © 2020 Michael R. Roberts

changes hands. This distinction is important for finance, which cares about when cash changes hands, not when they are recorded by accountants. To illustrate, consider two transactions and their cash flow implications.

1. We buy an Xbox 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 for a week or so from the credit card company.

2. We buy a three-year, $300 million contract for cloud computing services on Microsoft's Azure platform. 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 A.1 illustrates this distinction with a timeline. Regardless of when we pay, Microsoft recognizes the revenue according to the usage of the service.

Year	0	1	2	3
Sales ($mil)		100	100	100
Cash flows ($mil)		245		
Cash flows ($mil)				340

Figure A.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.

1.1.2 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(es), 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

Copyright © 2020 Michael R. Roberts

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

Copyright © 2020 Michael R. Roberts

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 an X-box 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 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 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. Let's consider an example.

Say 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

$$\text{Depreciation} = \frac{\text{Purchase price} - \text{Salvage value}}{\text{Length of usable life}} \tag{A.1}$$

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]

An important feature of depreciation and amortization to note is that they represent **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, we can think of a lot of instances where this is not true. So, you may ask, why bother recording this expense if it doesn't affect the financial value of the firm?

[1]To 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.

[2]Asset 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).

[3]In 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.

Copyright © 2020 Michael R. Roberts

What depreciation and amortization do is reduce a firm's taxable income and, by doing so, reduce the taxes it has to pay. In financial lingo, depreciation and amortization provide a *tax shield* for firms.These non-cash expenses reduce firm's tax burden.

1.1.3 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 and earnings 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 taxing 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

Copyright © 2020 Michael R. Roberts

bill from financial statements is difficult if not impossible. So, the tax expense item needs to be taken with a grain of salt.

1.1.4 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 deprecation and amortization are *non-cash* expenses.) It is also a measure of the cash earnings available for the firms creditors (i.e., people that lent money to the firm) *and* shareholders (i.e., owners). 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 (income) are by definition unrelated to the core operations of the company and are often largely influenced by the financial policy (i.e., 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.

Copyright © 2020 Michael R. Roberts

- 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 (or expense) measures exactly coincide with cash moving in or out of a company, though they can still be informative.

A.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)

Copyright © 2020 Michael R. Roberts

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. Like the P&L, it's useful to draw parallels between the corporate balance sheet and your personal balance sheet.

1.2.1 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 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

Copyright © 2020 Michael R. Roberts

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, net and gross. 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 impair $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.)

Copyright © 2020 Michael R. Roberts

1.2.2 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

[4]Accounts 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."

Copyright © 2020 Michael R. Roberts

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.

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 bares 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,

$$\text{Total assets} = \text{Total liabliities} + \text{Shareholders equity} \tag{A.2}$$

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.

[5]Limited liability prevents creditors from going after the personal assets of shareholders. So, the worst that can happy to a shareholder is that their stock becomes worthless.

Copyright © 2020 Michael R. Roberts

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.

A.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.

There are three sections to the statement showing cash flow from different parts of the company: operations, investing, and financing.

1.3.1 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

Copyright © 2020 Michael R. Roberts

	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	76,740
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	(27,577)
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	(48,515)
Net change in cash	$648

Table 5: Microsoft 2021 Statement of Cash Flows ($mil)

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,

Copyright © 2020 Michael R. Roberts

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.

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.

1.3.2 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.

[6] Accounts receivable and accounts payable are typically reported on a "net" basis, where net refers to net of bad debts.

Copyright © 2020 Michael R. Roberts

1.3.3 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. The 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.

1.3.4 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.

A.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.

Copyright © 2020 Michael R. Roberts

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.

1.4.1 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)**.

$$ROA_t = \frac{\text{Net income}_t}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{61,271}{(301,311 + 333,779)/2} = 0.193 \quad \text{(A.3)}$$

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.

$$ROA_t = \frac{\text{EBIT}_t \times (1-\tau)}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{69,916 \times (1-0.138)}{(301,311 + 333,779)/2} = 0.190 \quad \text{(A.4)}$$

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.

Copyright © 2020 Michael R. Roberts

2. **Return on investment (ROI)**. Also known as **Return on Invested Capital (ROIC)**.

$$\begin{aligned} ROI_t &= \frac{\text{EBIT}_t \times (1-\tau)}{((\text{Debt}_{t-1} + \text{Equity}_{t-1}) + (\text{Debt}_t + \text{Equity}_t)) \ / \ 2} \\ &= \frac{69,916 \times (1-0.138)}{((70,988 + 67775) + (118,304 + 141,988))/2} \\ &= 0.302 \end{aligned} \tag{A.5}$$

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)**.

$$ROE_t = \frac{\text{Net income}_t}{(\text{Equity}_{t-1} + \text{Equity}_t)/2} = \frac{61,271}{(118,340141,988)/2} = 0.471 \tag{A.6}$$

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 A.2 presents current and historical performance KPIs for Microsoft and two of its closest competitors - Alphabet (ticker=GOOG) and Amazon (ticker=AMZN). The figure show 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 later of which will be lower because of the lower risk faced by creditors.

Copyright © 2020 Michael R. Roberts

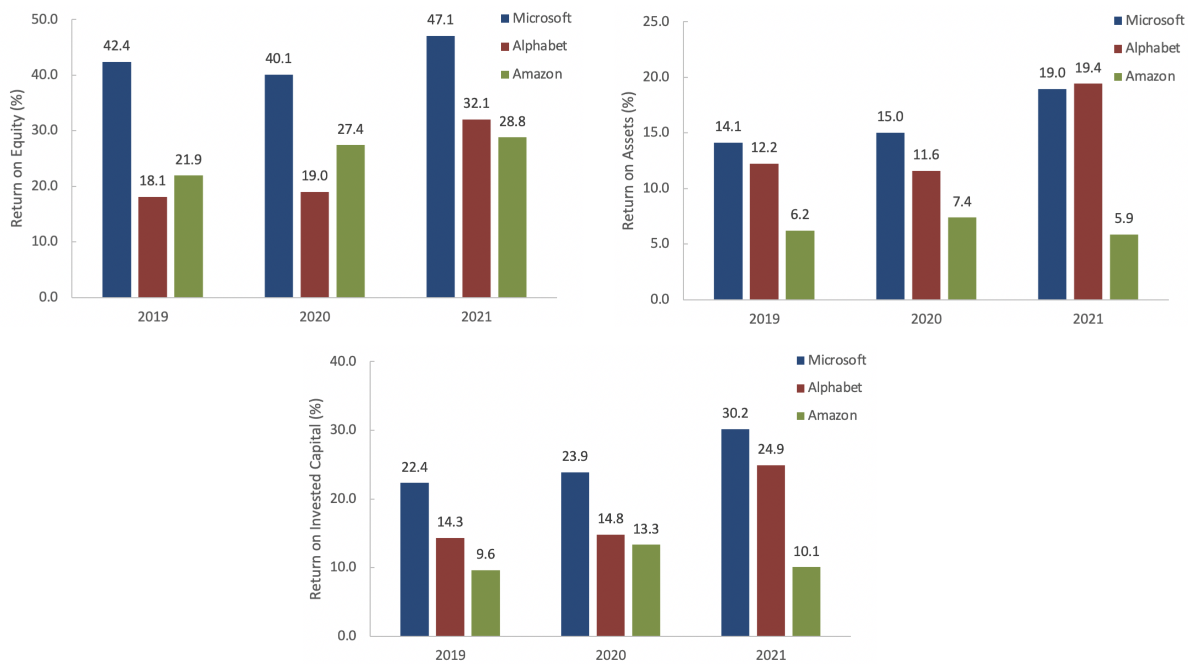

Figure A.2: Performance KPIs

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.*

Copyright © 2020 Michael R. Roberts

1.4.2 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.

	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 (%)

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.

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

Copyright © 2020 Michael R. Roberts

over the last three years, consistent with the decline in operating expense ratios. Figure A.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.

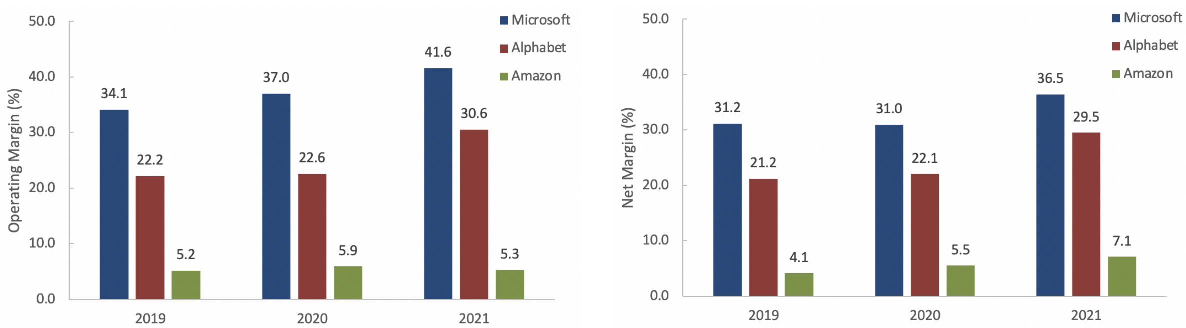

Figure A.3: Profit Margins

The expense ratio and margin trends show that Microsoft appears to be on an impressive trajectory. Before drawing any conclusions, it 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 A.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 A.2).

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.

1.4.3 Asset Efficiency

Asset turnover or **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

$$\text{Asset turnover}_t = \frac{\text{Sales}_t}{(\text{Assets}_{t-1} + \text{Assets}_t)/2} = \frac{168,088}{(301,311 + 333,779)/2} = 0.5$$

Copyright © 2020 Michael R. Roberts

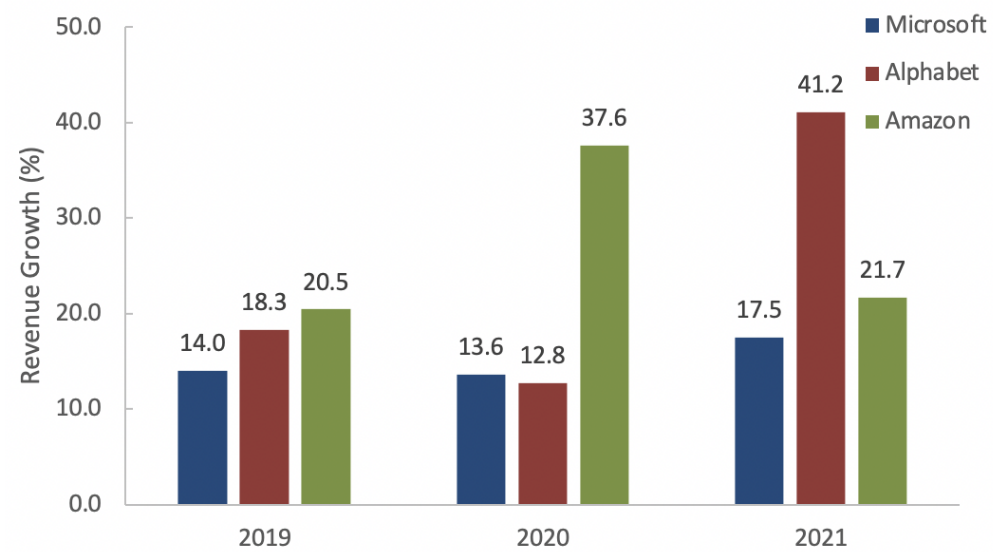

Figure A.4: Revenue Growth

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 A.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

Copyright © 2020 Michael R. Roberts

models to a large degree. Microsoft and Amazon have significant retail arms, whereas Alphabet is almost exclusively software and services.

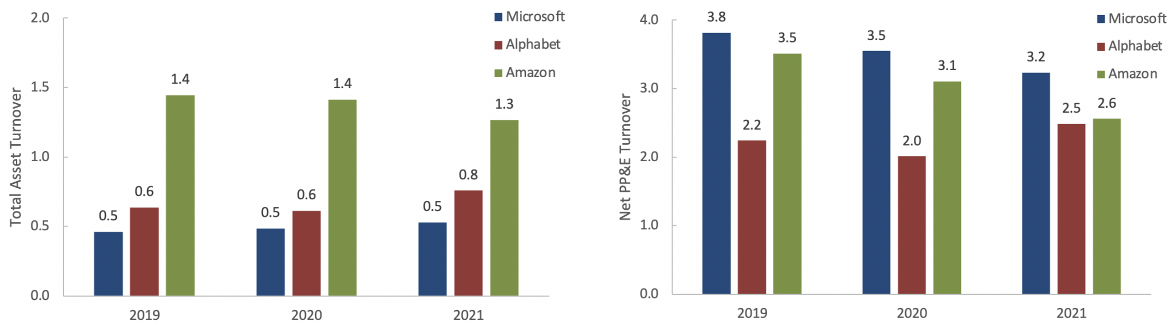

Figure A.5: Asset Turnover/Utilization Ratios

1.4.4 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**.

$$\text{Current ratio}_t = \frac{\text{Current assets}_t}{\text{Current liabilities}_t} = \frac{184,406}{88,657} = 2.1 \tag{A.7}$$

   Microsoft has$2.10 of liquid assets for every dollar that it owes over the next year.

2. **Quick ratio**.

$$\text{Current ratio}_t = \frac{\text{Cash}_t + \text{Short-term investments}_t + \text{Receivables}_t}{\text{Current liabilities}_t} = \frac{168,377}{88,657} = 1.9 \tag{A.8}$$

   (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 A.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

Copyright © 2020 Michael R. Roberts

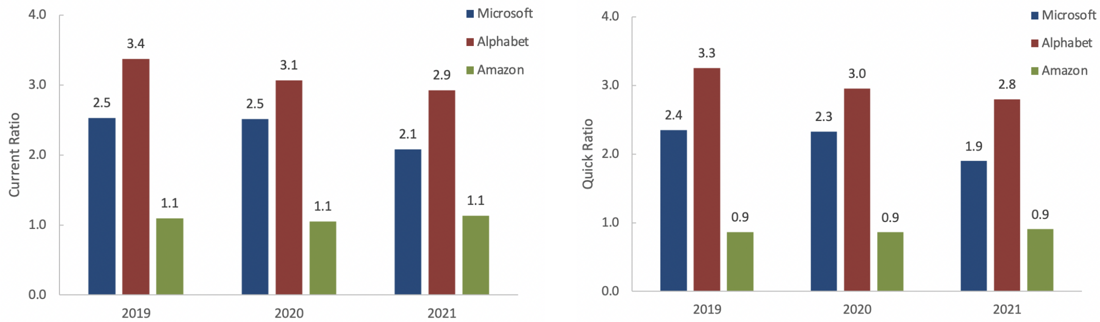

Figure A.6: Liquidity Measures

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. Its 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.

1.4.5 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 re-**

[7] Often 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.

Copyright © 2020 Michael R. Roberts

ceivable days and **accounts receivable collection period**.

$$\begin{aligned} DSO_t &= \frac{(\text{Receivables}_{t-1} + \text{Receivables}_t)/2}{\text{Sales}_t} \times 365 \\ &= \frac{(32,011 + 38,043)/2}{168,088} \times 365 \\ &= 76.1 \end{aligned} \tag{A.9}$$

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 A.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 A.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.)

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 A.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. It's end-of-year receivables will equal its sales for the year. Equation A.9 would imply that its DSO is equal to 365 when in fact it may require far fewer days

Copyright © 2020 Michael R. Roberts

to collect. Lesson: Be careful using equation A.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**.

$$\begin{aligned} DPO_t &= \frac{(\text{Payables}_{t-1} + \text{Payables}_t)/2}{\text{Cost of sales}_t} \times 365 \qquad (A.10) \\ &= \frac{(12,530 + 15,163)/2}{40,546} \times 365 \\ &= 124.6 \end{aligned}$$

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 A.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**.

$$\begin{aligned} DIO_t &= \frac{(\text{Inventory}_{t-1} + \text{Inventory}_t)/2}{\text{Cost of sales}_t} \times 365 \qquad (A.11) \\ &= \frac{(1,895 + 2,636)/2}{40,546} \times 365 \\ &= 20.4 \end{aligned}$$

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**.

$$CCC_t = DSO_t + DIO_t - DPO_t = 76.1 + 20.4 - 124.6 = -28.2 \qquad (A.12)$$

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

Copyright © 2020 Michael R. Roberts

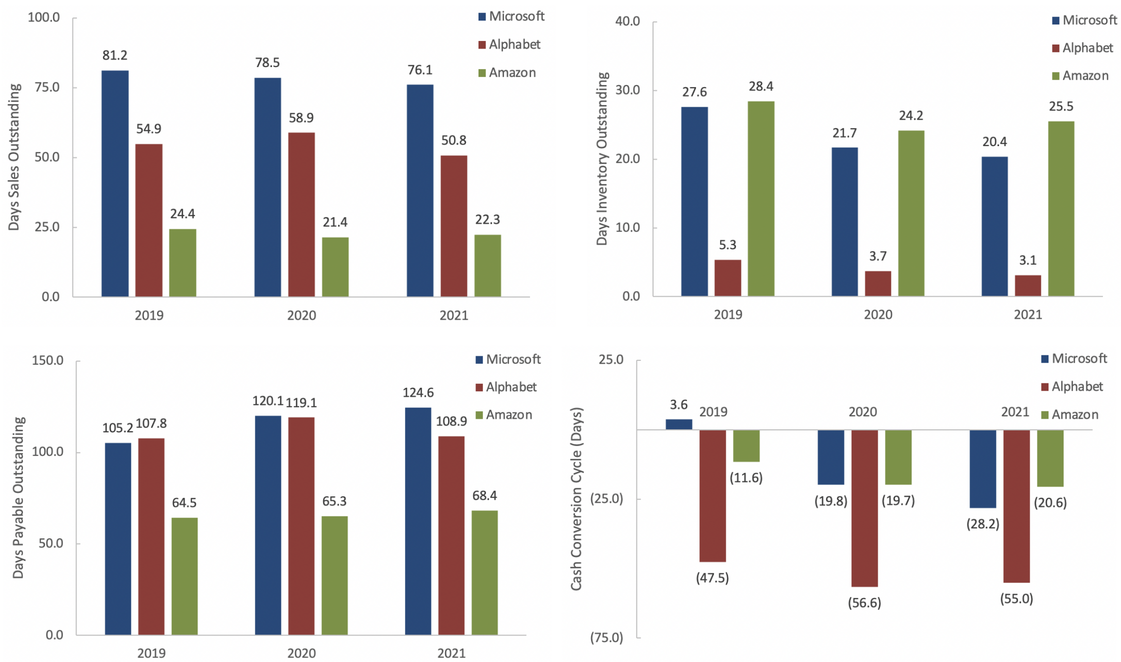

Figure A.7: Working Capital KPIs

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 A.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

Copyright © 2020 Michael R. Roberts

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.

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.

1.4.6 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**.

$$\text{Debt-to-ebitda}_t = \frac{\text{Total debt}_t}{\text{EBITDA}_t} = \frac{67,775}{81,602} = 0.8\mathbf{x} \tag{A.13}$$

   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**.

$$\text{Interest coverage}_t = \frac{\text{EBITDA}_t}{\text{Interest expense}_t} = \frac{81,602}{-1,186} = -59.0\mathbf{x} \tag{A.14}$$

[8] As 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

Copyright © 2020 Michael R. Roberts

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**.

$$\begin{aligned} \text{Debt service coverage}_t &= \frac{\text{EBITDA}_t}{\text{Interest expense}_t + \text{Debt due this year}_t} \quad \text{(A.15)} \\ &= \frac{81,602}{-1,186 + 8,072} \\ &= 11.9\mathbf{x} \end{aligned}$$

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**.

$$\begin{aligned} \text{Debt-to-total capitalization}_t &= \frac{\text{Total debt}_t}{\text{Total debt}_t + \text{Book equity}_t} \quad \text{(A.16)} \\ &= \frac{67,775}{67,775 + 141,988} \\ &= 0.323 \quad \text{(A.17)} \end{aligned}$$

The debt-to-total capitalization ratio measures the fraction of the firm's financing to came 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**.

$$\begin{aligned} \text{Net debt-to-total capitalization}_t &= \frac{\text{Total debt}_t - \text{Cash}_t}{\text{Total debt}_t - \text{Cash}_t + \text{Book equity}_t} \quad \text{(A.18)} \\ &= \frac{67,775 - 130,334}{67,775 - 130,334 + 141,988} \\ &= 0.788 \quad \text{(A.19)} \end{aligned}$$

Cash in this formula refers to cash textitand short-term investments. The net debt-to-total capitalization ratio recognizes that firms can often use their cash holdings to pay

Copyright © 2020 Michael R. Roberts

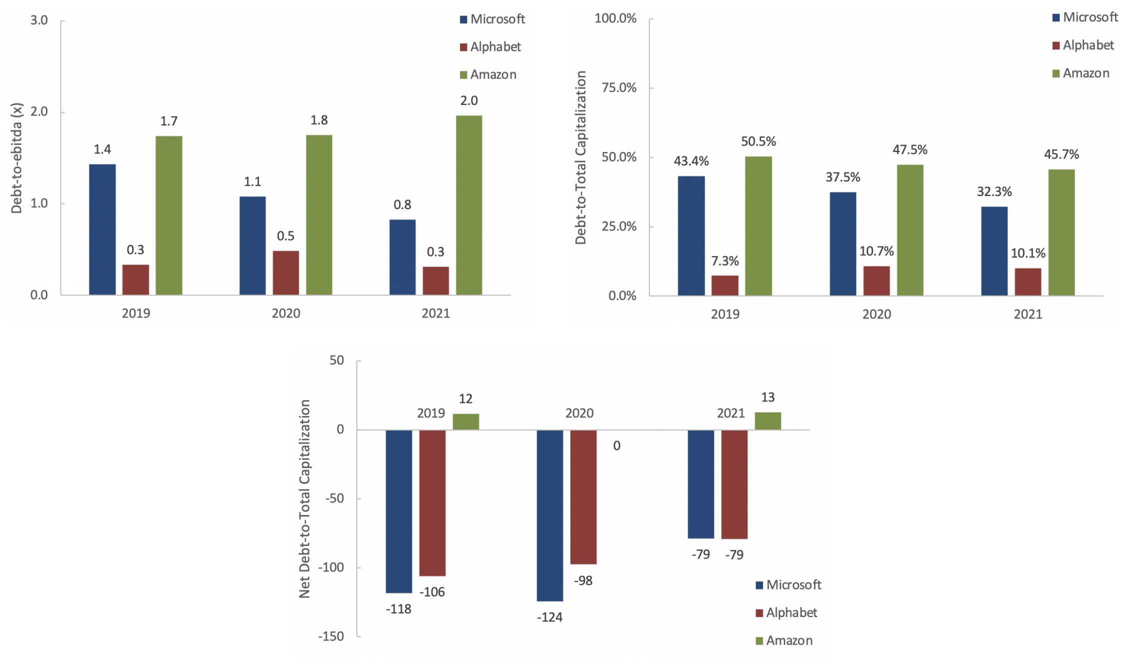

Figure A.8: Credit Risk KPIs

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 A.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

Copyright © 2020 Michael R. Roberts

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.

1.4.7 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]

$$\underbrace{\frac{\text{Net income}}{\text{Equity}}}_{ROE} = \overbrace{\underbrace{\frac{\text{Net income}}{\text{Sales}}}_{\text{Net margin}} \times \underbrace{\frac{\text{Sales}}{\text{Assets}}}_{\text{Asset turnover/utilization}}}^{ROA} \times \underbrace{\frac{\text{Assets}}{\text{Equity}}}_{\text{Financial leverage}} \tag{A.20}$$

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 A.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.

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.

[9]The name comes from the DuPont Corporation, the company at which it was invented.

Copyright © 2020 Michael R. Roberts

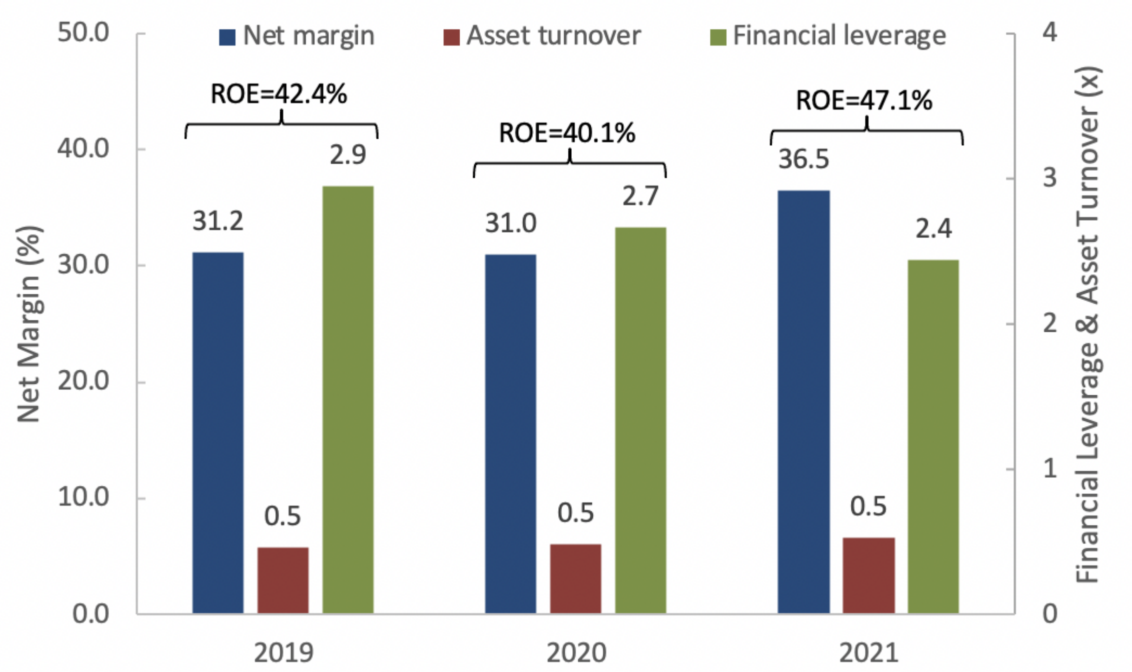

Figure A.9: Du Pont Analysis of Microsoft's ROE

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 a a very modest rate of 4.5% and depreciation and amortization contracted by 8.5%. We also see a very large increase 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 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.

1.4.8 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)

Copyright © 2020 Michael R. Roberts

	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

space are particularly found 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.

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 A.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

Copyright © 2020 Michael R. Roberts

	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

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.

A.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

Copyright © 2020 Michael R. Roberts

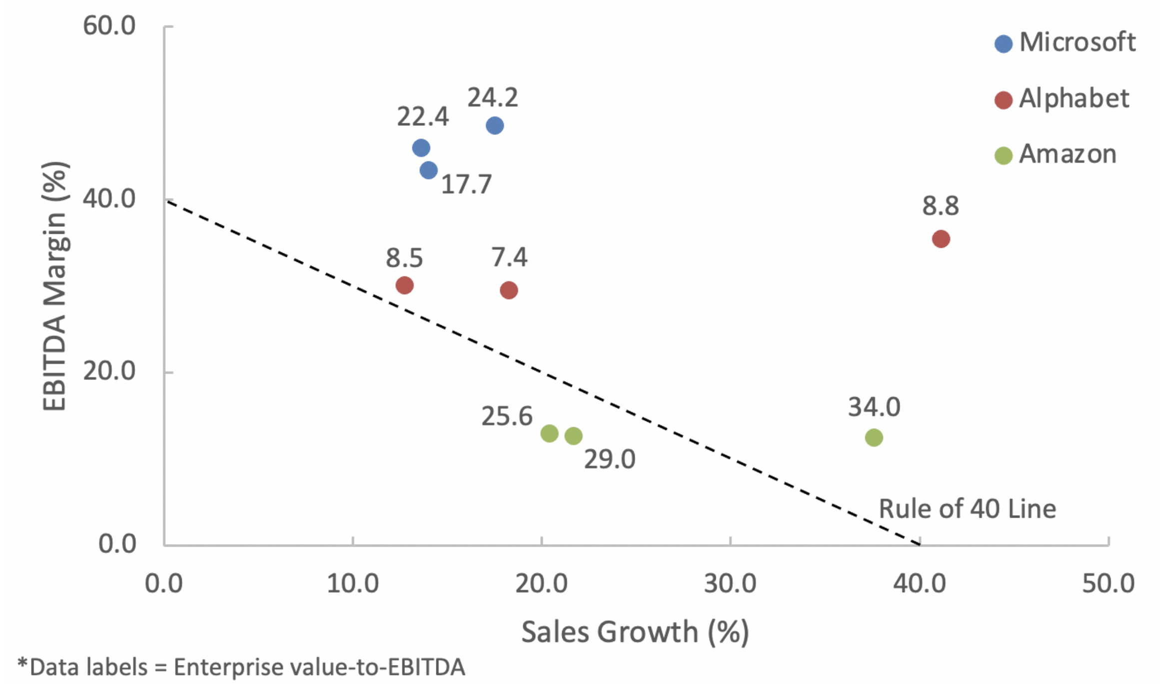

Figure A.10: Rule of 40 Analysis

(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.

Copyright © 2020 Michael R. Roberts

Appendix B

References