Risk, Cost of Capital, and Valuation 12

OPENING CASE

With over 110,000 employees on five continents, Germany-based BASF is a major international company. It operates in a variety of industries, including agriculture, oil and gas, chemicals, and plastics. In an attempt to increase value, BASF launched Vision 2020, a comprehensive plan that included all functions within the company and challenged and encouraged all employees to act in an entrepreneurial manner. The major financial component of the strategy was that the company expected to earn its weighted average cost of capital, or WACC, plus a premium. So, what exactly is the WACC?

The WACC is the minimum return a company needs to earn to satisfy all of its investors, including stockholders, bondholders, and preferred stockholders. In 2011, for example, BASF pegged its cost of capital at 11 percent and earned a premium of €2.55 billion above its cost of capital. From 2012 to 2015, the company kept its WACC pegged at 11 percent. In this chapter, we learn how to compute a firm's cost of capital and find out what it means to the firm and its investors. We will also learn when to use the firm's cost of capital and, perhaps more important, when not to use it.

Please visit us at corecorporatefinance.blogspot.com for the latest developments in the world of corporate finance.

The goal of this chapter is to determine the rate at which cash flows of risky projects are to be discounted. Projects are financed with equity, debt, and other sources, and we must estimate the cost of each of these sources in order to determine the appropriate discount rate. We begin with the cost of equity capital. Since the analysis here builds on beta and the capital asset pricing model (CAPM), we discuss beta in depth, including its calculation, its intuition, and its determinants. We next discuss the cost of debt and the cost of preferred stock. These costs serve as building blocks for the weighted average cost of capital (R_{WACC} or, WACC), which is used to discount cash flows. We calculate the WACC for a real-world company, Eastman Chemical Co. Finally, we introduce flotation costs.

12.1 THE COST OF EQUITY CAPITAL

Whenever a firm has extra cash, it can take one of two actions. It can pay out the cash directly to its investors. Alternatively, the firm can invest the extra cash in a project, paying out the future cash flows of the project. Which action would the investors prefer? If an investor can reinvest the cash in a financial asset (a stock or bond) with the same risk as that of the project, the investors would desire the alternative with the highest expected return. In other words, the project should be undertaken only if its expected return is greater

than that of a financial asset of comparable risk. This idea is illustrated in Figure 12.1. Our discussion implies a very simple capital budgeting rule: page 358

The discount rate of a project should be the expected return on a financial asset of comparable risk.

FIGURE 12.1
Choices of a Firm with Extra Cash

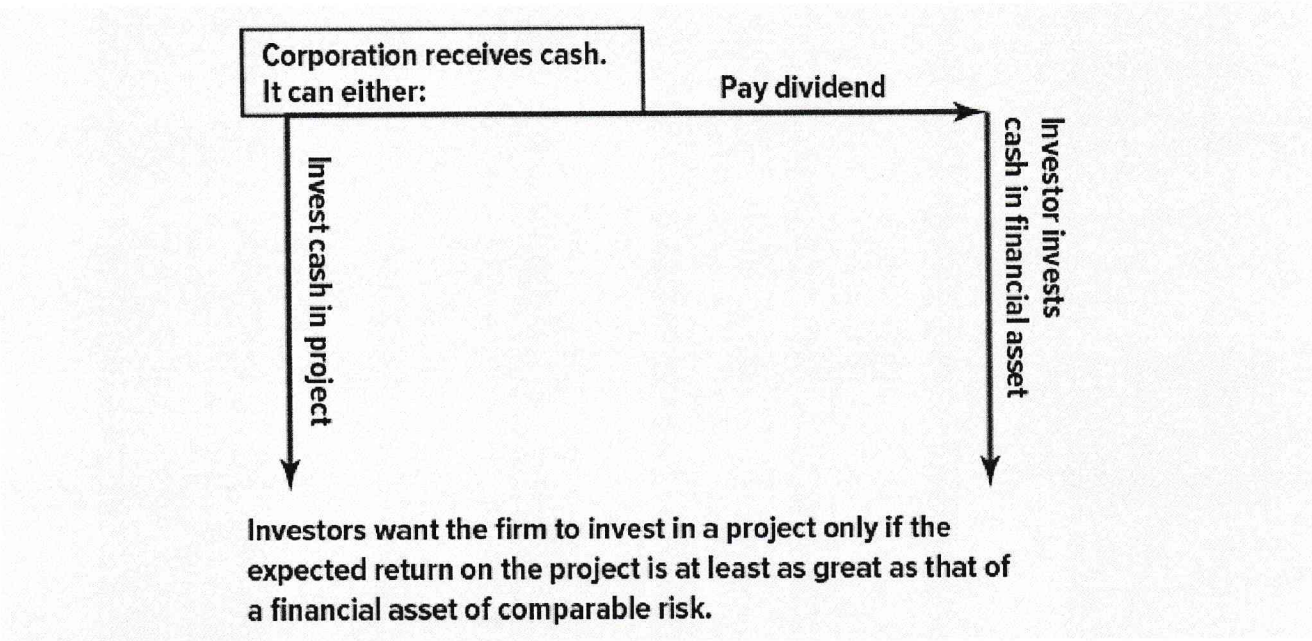

There are various synonyms for the discount rate. For example, the discount rate is often called the *required return* on the project. This is an appropriate name, since the project should be accepted only if the project generates a return above what is required. Alternatively, the discount rate of the project is said to be its *cost of capital.* This name is also appropriate, since the project must earn enough to pay its suppliers of capital. Our book will use these three terms: the discount rate, the required return, and the cost of capital, synonymously.

Now imagine that all projects of the firm have the same risk. In that case, one could say that the discount rate is equal to the cost of capital for the firm as a whole. And, if the firm is all equity, the discount rate is also equal to the firm's cost of equity capital.

12.2 ESTIMATING THE COST OF EQUITY CAPITAL WITH THE CAPM

It's one thing to define the cost of equity capital, as we have done above. It's quite another to estimate it. The problem is that stockholders do not tell the firm what their required returns are. So, what do we do? Luckily, the capital asset pricing model (CAPM) can be used to estimate the required return.

Under the CAPM, the expected return on the stock can be written as:

$$R_S = R_F + \beta \times (R_M - R_F) \qquad [12.1]$$

where R_F is the risk-free rate and $R_M - R_F$ is the difference between the expected return on the market portfolio and the riskless rate. Recall that this difference is often called the expected *excess* market return or market risk premium. Note we have dropped the bar denoting expectations from our expression to simplify the notation, but remember that we are always thinking about *expected* returns with the CAPM.

The expected return on the stock in Equation 12.1 is based on the stock's risk, as measured by beta. Alternatively, we could say that this expected return is the required return on the stock, based on the stock's risk. Similarly, this expected return can be viewed as the firm's cost of equity capital.

Market Risk Premium

METHOD 1: USING HISTORICAL DATA Much of Chapter 10 was devoted to the calculation of historical rates of return and the market risk premium. The chapter settled on an estimate of 7 percent for the premium, though this number should not be interpreted as definitive. Next we need a risk-free rate. The current one-year Treasury bill is about .70 percent.

As a quick example, consider an all-equity company with a beta of 1.5. Given our parameters, its cost of capital would be:

$$.70\% + 1.5 \times 7\% = 11.20\%$$

METHOD 2: USING THE DIVIDEND DISCOUNT MODEL (DDM) Earlier in this chapter, we referenced a study indicating that most corporations use the CAPM for capital budgeting. Does the CAPM imply that risk premiums must be calculated from past returns, as we did above? The answer is no. There is another method, based on the dividend discount model of an earlier chapter, for estimating the risk premium.

In Chapter 6, we pointed out that the price of a share of stock is equal to the present value of all of its future dividends. Furthermore, we noted in that chapter that, if the firm's dividends are expected to grow at a constant rate, g, the price of a share of stock, P, can be written as:

$$P = \frac{\text{Div}}{R - g}$$

where Div is the dividend per share to be received next year, R is the discount rate, and g is the constant annual rate of growth in dividends. This equation can be rearranged, yielding:

$$R = \frac{\text{Div}}{P} + g$$

In words, the annual return on a stock is the sum of the dividend yield (=Div/P) over the next year plus the annual growth rate in dividends.

Just as this formula can be used to estimate the total return on a stock, it can be used to estimate the total return on the market as a whole. The first term on the right-hand side is easy to estimate, since a number of print and web services calculate the dividend yield for the market. For example, as this is written, the average dividend yield across all stocks in the Standard & Poor's (S&P) 500 Index was about 2.2 percent. We will use this number in our forecasts.

Next, we need an estimate of the per-share growth rate in dividends across all companies in the market. Security analysts, who are typically employees of investment banking houses, money management firms, and independent research organizations, study individual securities, industries, and the overall stock market. As part of their work, they forecast dividends and earnings, as well as make stock recommendations. For example, suppose the numbers in the Value Line (VL) *Investment Survey* imply a five-year growth rate in dividends for VL's Industrial Composite Index of about 6 percent per

year. With a dividend yield of 2.2 percent, the expected return on the market becomes 2.2% + 6% = 8.2%. Given our one-year yield on Treasury bills of .70 percent, the market risk premium would be 8.2% − .70% = 7.50%, a number somewhat above the 7 percent provided by Method 1.

For our firm with a beta of 1.5, the cost of capital becomes:

$$.70\% + 1.5 \times 7.50\% = 11.95\%$$

Of course, Value Line is just one source for forecasts. More likely, a firm would either rely on a consensus of many forecasts or use its own subjective growth estimate.

Academics have, nevertheless, long preferred the historical market risk premium for its objectivity. Since historical returns have been precisely measured, there is little room

for subjective judgment. By contrast, estimation of future dividend growth in the DDM is more subjective. However, the subjective nature of the DDM approach is not meant as a criticism. Proponents of using the DDM point out that returns in the long run can only come from the current dividend yield and future dividend growth. Anyone who thinks that long-run stock returns will exceed the sum of these two components is fooling himself.[3] The expression, "You can't squeeze blood out of a turnip," applies here.

page 362

12.3 ESTIMATION OF BETA

In the previous section, we assumed that the beta of the company was known. Of course, beta must be estimated in the real world. We pointed out earlier that the beta of a security is the standardized covariance of a security's return with the return on the market portfolio. As we have seen, the formula for Security i is:

$$\text{Beta of Security } i = \frac{\text{Cov}(R_i, R_M)}{\text{Var}(R_M)} = \frac{\sigma_{i,M}}{\sigma_M^2} \qquad [12.2]$$

In words, the beta is the covariance of a security with the market, divided by the variance of the market. Because we calculated both covariance and variance in earlier chapters, calculating beta involves no new material.

Measuring Company Betas

The basic method of measuring company betas is to estimate:

$$\frac{\text{Cov}(R_i, R_M)}{(\text{Var}R_M)}$$

using $t = 1, 2, \ldots, T$ observations.

Problems

1. Betas may vary over time.
2. The sample size may be inadequate.
3. Betas are influenced by changing financial leverage and business risk.

Solutions

1. Problems 1 and 2 can be moderated by more sophisticated statistical techniques.
2. Problem 3 can be lessened by adjusting for changes in business and financial risk.
3. Look at average beta estimates of several comparable firms in the industry.

Real-World Betas

It is instructive to see how betas are determined for actual real-world companies. Figure 12.3 plots monthly returns for four large firms against monthly returns on the Standard & Poor's (S&P) 500 Index. Using a standard regression technique, we fit a straight line through the data points. The result is called the "characteristic" line for the security. The slope of the characteristic line is beta. Though we have not shown it in the figure, we can also determine the intercept (commonly called alpha) of the characteristic line by regression.

We use five years of monthly data for each plot. Although this choice is arbitrary, it is in line with calculations performed in the real world. Practitioners know that the accuracy

TABLE 12.1 Betas for Firms in the Computer Software Industry

COMPANY	BETA
Microsoft	1.04
Apple, Inc.	1.36
Automatic Data Processing	.84
Oracle Corp.	1.03
Computer Sciences	.71
CA, Inc.	.60
Fiserv, Inc.	.78
Accenture, Ltd.	.93
Symantec Corp.	1.01
Paychex, Inc.	.98
Equally-weighted portfolio	.93

Assuming a risk-free rate of .70 percent and a risk premium of 7 percent, Automatic Data Processing might estimate its cost of equity capital as

$$.70\% + .84 \times 7\% = 6.58\%$$

However, if Automatic Data Processing believed the industry beta contained less estimation error, it could estimate its cost of equity capital as

$$.70\% + .93 \times 7\% = 7.20\%$$

The difference is relatively small here, so it is a simple choice for a financial executive at Automatic Data Processing. However, the difference can often be much larger, which presents a much more difficult choice for financial executives.

While there is no formula for selecting the right beta, there is a very simple guideline. If you believe that the operations of a firm are similar to the operations of the rest of the industry, you should use the industry beta to reduce estimation error.[6] However, if an executive believes that the operations of the firm are fundamentally different from those in the rest of the industry, the firm's beta should be used.

When we discussed financial statement analysis in Chapter 3, we noted that a problem frequently comes up in practice—namely, what is the industry? For example, Value Line's *Investment Survey* categorizes Accenture, Ltd., as a computer software company, whereas online financial providers such

as www.reuters.com/finance categorize the same company in the business services industry, so more than a little care must be taken in using the industry approach.

12.4 DETERMINANTS OF BETA

The regression analysis approach in Section 12.3 doesn't tell us where beta comes from. Of course, the beta of a stock does not come out of thin air. Rather, it is determined by the characteristics of the firm. We consider three factors: the cyclical nature of revenues, operating leverage, and financial leverage.

Cyclicality of Revenues

The revenues of some firms are quite cyclical. That is, these firms do well in the expansion phase of the business cycle and do poorly in the contraction phase. Empirical evidence suggests high-tech firms, retailers, and automotive firms fluctuate with the business cycle. Firms in industries such as utilities, railroads, food, and airlines are less dependent on the cycle. Because beta measures the responsiveness of a stock's return to the market's return, it is not surprising that highly cyclical stocks have high betas.

It is worthwhile to point out that cyclicality is not the same as variability. For example, a moviemaking firm has highly variable revenues because hits and flops are not easily predicted. However, because the revenues of a studio are more dependent on the quality of its releases than the phase of the business cycle, motion picture companies are not particularly cyclical. In other words, stocks with high standard deviations need not have high betas, a point we have stressed before.

Operating Leverage

We distinguished fixed costs from variable costs in Chapter 9. At that time, we mentioned that fixed costs do not change as quantity changes. Conversely, variable costs increase as the quantity of output rises. Firms often face a trade-off between fixed and variable costs. For example, a firm can build its own factory, incurring a high level of fixed costs in the process. Alternatively, the firm can outsource production to a supplier, typically generating lower fixed costs but higher variable costs. Fixed costs tend to magnify the impact of sales cyclicality. Fixed costs must be paid, even at a low level of sales, leaving the firm with the possibility of large losses. And with fixed costs replacing variable costs, any additional sales generate low marginal costs, leaving the firm with a substantial increase in profit.

Firms with high fixed costs and low variable costs are generally said to have high **operating leverage**. Conversely, firms with low fixed and high variable costs have low operating leverage. Operating leverage magnifies the effect of the cyclicality of a firm's revenues on beta. That is, a firm with a given sales cyclicality will increase its beta if fixed costs replace variable costs in its production process.

Financial Leverage and Beta

As suggested by their names, operating leverage and financial leverage are analogous concepts. Operating leverage refers to the firm's fixed costs of *production.* Financial leverage is the extent to which a firm relies on debt, and a levered firm is a firm with some debt in its capital structure. Because a *levered* firm must make interest payments regardless of the firm's sales, financial leverage refers to the firm's fixed costs of *finance.*

Just as an increase in operating leverage increases beta, an increase in financial leverage (i.e., an increase in debt) increases beta. To see this point, consider a firm with some debt and some equity in its capital structure. Further, imagine an individual who owns all the firm's debt and all its equity. In other words, this individual owns the entire firm. What is the beta of her portfolio of the firm's debt and equity?

As with any portfolio, the beta of this portfolio is a weighted average of the betas of the individual items in the portfolio. Let B stand for the market value of the firm's debt and S stand for the market value of the firm's equity. We have:

$$\beta_{Portfolio} = \beta_{Asset} = \frac{S}{B+S} \times \beta_{Equity} + \frac{B}{B+S} \times \beta_{Debt} \qquad [12.3]$$

where β_{Equity} is the beta of the stock of the *levered* firm. Notice that the beta of debt, (β_{Debt}, is multiplied by $B/(B + S)$, the percentage of debt in the capital structure. Similarly, the beta of equity is

multiplied by the percentage of equity in the capital structure. Because the portfolio contains both the debt of the firm and the equity of the firm, the beta of the portfolio can be thought of as the beta of the common stock had the firm been all equity. In practice, this beta is called the **asset beta** because its value is dependent only on the assets of the firm.

The beta of debt is very low in practice. If we make the common assumption that the beta of debt is zero, we have:

$$\beta_{Asset} = \frac{S}{B+S} \times \beta_{Equity} \qquad [12.4]$$

12.6 COST OF CAPITAL FOR DIVISIONS AND PROJECTS

Previous sections of this chapter all assumed that the risk of a potential project is equal to the risk of the existing firm. How should we estimate the discount rate for a project whose risk differs from that of the firm? The answer is that each project should be discounted at a rate commensurate with its own risk. For example, let's assume that we use the CAPM to determine the discount rate.[10] If a project's beta differs from that of the firm, the project's cash flows should be discounted at a rate commensurate with the project's own beta. This is an important point, since firms frequently speak of a *corporate discount rate.* (As mentioned earlier, *required return* and *cost of capital* are frequently used synonymously.) Unless all projects in the corporation are of the same risk, choosing the same discount rate for all projects is incorrect.

The above paragraph considered the discount rates of individual projects. The same message would apply for whole divisions. If a corporation has a number of divisions, each in a different industry, it would be a mistake to assign the same discount rate to each division.

EXAMPLE 12.4

Project Risk

D. D. Ronnelley Co., a publishing firm, may accept a project in computer software. Noting that computer software companies have high betas, the publishing firm views the software venture as more risky than the rest of its business. It should discount the project at a rate commensurate with the risk of software companies. For example, it might use the average beta of a portfolio of publicly traded software firms. Instead, if all projects in D. D. Ronnelley Co. were discounted at the same rate, a bias would result. The firm would accept too many high-risk projects (software ventures) and reject too many low-risk projects (books and magazines). This point is illustrated in Figure 12.5.

FIGURE 12.5
Relationship between the Firm's Cost of Capital and the Security Market Line (SML)

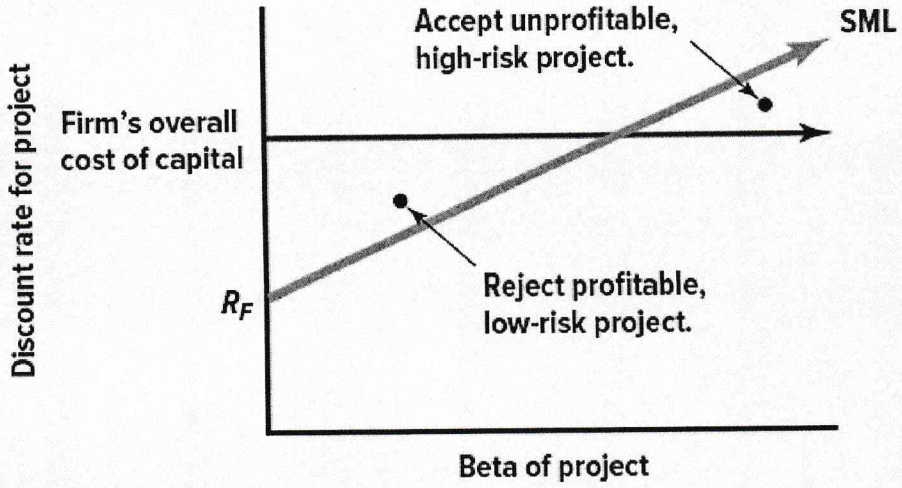

A single cost of capital for all projects in a firm, as indicated by the horizontal line in the figure, may lead to incorrect capital budgeting decisions. Projects with high risk, such as the software venture for D. D. Ronnelley Co., should be discounted at a high rate. By using the firm's cost of capital, the firm is likely to accept too many high-risk projects.

Projects with low risk should be discounted at a low rate. By using the firm's cost of capital, the firm is likely to reject too many low-risk projects.

The D. D. Ronnelley (DDR) example points out that we should discount a project at a rate commensurate with the risk of the project's cash flows. However, practitioners should be concerned with three issues here. First, they must choose the appropriate industry. While this may seem to be an easy task, the problem is that companies often have more than one line of business. For example, suppose that DDR was considering a project in the movie industry, not in computer software. Their first thought might be to look at the betas of the largest and most important companies in the film industry. The six biggest studios are Warner Brothers, Columbia, Universal, Paramount, 21st Century Fox, and Disney. However, the first four studios are owned by Time-Warner, Sony, News Corporation, Comcast, and Viacom, respectively. These parent corporations are all diversified, with movies making up only a small portion of total revenues. 21st Century Fox does produce movies, but it is also heavily involved in news, sports, cable television and network ownership. And while the parent of the sixth studio has the same Walt Disney name, it too is quite diversified, with holdings in television, radio, theme parks, and cruise ships. With all this diversification, it would likely be quite difficult to determine the beta of a pure moviemaking company from the betas of the six parents. Analysts often talk about identifying *pure plays* (i.e., other companies that specialize only in projects similar to the project your firm is considering). Pure plays are easier to find in some situations than in others.

Second, even if all companies in a particular industry are pure plays, the beta of a new project may be greater than the beta of existing firms, because a new project is likely to be particularly responsive to economywide movements. For example, a start-up computer venture may fail in a recession while IBM, Microsoft, or Oracle will still be around. Conversely, in an expansion, the venture may grow faster than the older computer firms.

Fortunately, a slight adjustment is all that is needed here. The new venture should be assigned a somewhat higher beta than that of the industry to reflect added risk. The adjustment is necessarily ad hoc, so no formula can be given. Our experience indicates that this approach is in widespread practice today.

Third, a problem arises for the rare project constituting its own industry. For example, consider the firms providing consumer shopping by television. Today, we can obtain a reasonable estimate for the beta of this industry because a few of the firms have publicly traded stock. However, when the ventures began in the 1980s, any beta estimate was suspect. At that time, no one knew whether shopping by TV belonged in the television industry, the retail industry, or in an entirely new industry.

What beta should be used when the project constitutes its own industry? Earlier in this chapter we mentioned three determinants of beta: Cyclicality of revenues, operating leverage, and financial leverage. Comparing the values of these three determinants for the project in question to the values for other firms should provide at least a general feel for the project's beta.

12.7 COST OF FIXED INCOME SECURITIES

In this section, we examine the cost of both debt and preferred stock. We consider the cost of debt first.

Cost of Debt

The cost of equity is often difficult to estimate. The task generally involves a fair amount of data gathering, and the end result is often measured with error. Fortunately, the cost of debt is much easier to determine; it is the cost of borrowing. The firm can generally obtain this information either by checking the yield on publicly traded bonds or by talking to commercial and investment bankers.

Two years ago, the Ritter Manufacturing Corp. (RMC) issued $100 million of debt with a 7 percent coupon. While the bonds were initially issued at par, rising interest rates over the last two years have caused them to sell at a discount. The yield on the bonds is currently 8 percent. In order to finance expansion, RMC is considering another large issue of bonds. What is the cost of the new debt?

The cost of the new debt should be around 8 percent. If the old bonds are selling at 8 percent, the new debt will not sell at a lower yield. The 7 percent is merely a historical number, often called the *embedded cost* of the debt, with no relevance today.

Alternatively, perhaps a firm is issuing debt for the first time. Here, the firm's investment banker can generally indicate to the firm's managers what the yield on the prospective bonds will be. That yield can be used as an estimate of the cost of debt.[11] Or, perhaps the company will take out a loan with a commercial bank. Again, the borrowing rate on the prospective loan is the cost of debt.

There is only one complication that needs to be discussed. We have ignored taxes so far, obviously an assumption at odds with reality. Under U.S. tax law, interest payments are *tax deductible.* Consider the following example where two firms, Unlevered Corp. and Levered Corp., differ only in debt. Unlevered Corp. has no debt, and Levered Corp. has $100 of debt with an interest rate of 10 percent.

UNLEVERED CORP.	
Revenue	$180
Expenses	−70
Pretax earnings	110
Taxes (40% rate)	−44
Aftertax earnings	$ 66

LEVERED CORP.	
Revenue	$180
Expenses	−70
Earnings before interest and taxes	110
Interest (10% on $100 borrowed)	−10
Pretax earnings	100
Taxes (40% rate)	−40
Aftertax earnings	$ 60

While the Levered Corp. must pay $10 of interest per year, its aftertax earnings are only $6 (= $66 − 60) less than those of the Unlevered Corp. Why? Because the interest payments are tax deductible. That is, while Levered Corp.'s pretax earnings are $10 (= $110 − 100) less than those of Unlevered Corp., Levered Corp. pays $4 (= $44 − 40) less in taxes than does Unlevered Corp.

The $6 reduction of aftertax earnings is 6 percent of the $100 that Levered Corp. borrowed. Thus, the aftertax cost of debt is 6 percent. In general, the aftertax cost of debt can be written as:

$$\text{Aftertax cost of debt} = (1 - \text{Tax rate}) \times \text{Borrowing rate}$$

$$6\% = (1 - .40) \times 10\%$$

Why have we tax-adjusted the cost of debt but not the cost of equity? Because, while firms can deduct their interest payments before paying taxes, dividends are not tax deductible.

Cost of Preferred Stock

The name preferred stock is an unfortunate one, because preferred stock is probably more similar to bonds than to common stock. Preferred stock pays a constant dividend in perpetuity. Interest payments on bonds are quite similar to dividends on preferred stock,

though almost all bonds have a finite maturity. By contrast, dividends on common stock are not constant over time. page 372

Suppose a share of the preferred stock of Polytech, Inc., is selling at $17.16 and pays a dividend of $1.50 per year. Since preferred stocks are perpetuities, they should be priced using the perpetuity formula, PV = C/R, where PV is the present value, or price, C is the cash to be received each year, and r is the yield, or rate of return. Rearranging, we have:

$$R = C/PV$$

For this preferred issue, the rate of return is 8.7% (= $1.50/17.16). The cost of preferred stock is this rate of return.

Why don't we tax-adjust the cost of preferred stock the way we did the cost of debt? We don't tax-adjust here because dividend payments on preferred stock are not tax deductible.

12.8 THE WEIGHTED AVERAGE COST OF CAPITAL

Sections 12.1 and 12.2 showed how to estimate the discount rate when a project is all equity financed. In this section, we discuss an adjustment when the project is financed with both debt and equity.

Suppose a firm uses both debt and equity to finance its investments. If the firm pays R_B for its debt financing and R_S for its equity, what is the overall or average cost of its capital? The cost of equity is R_S, as discussed in earlier sections. The cost of debt is the firm's borrowing rate, R_B, which we can often observe by looking at the yield to maturity on the firm's debt. If a firm uses both debt and equity, the cost of capital is a weighted average of each. This works out to be:

$$\frac{S}{S+B} \times R_S + \frac{B}{S+B} \times R_B$$

The weights in the formula are, respectively, the proportion of total value represented by equity:

$$\left(\frac{S}{S+B}\right)$$

and the proportion of total value represented by debt:

$$\left(\frac{B}{S+B}\right)$$

This is only natural. If the firm had issued no debt and was therefore an all-equity firm, its average cost of capital would equal its cost of equity, R_S. At the other extreme, if the firm had issued so much debt that its equity was valueless, it would be an all-debt firm, and its average cost of capital would be its cost of debt, R_B.

Interest is tax deductible at the corporate level, as stated in the previous section. The aftertax cost of debt is:

$$\text{Cost of debt (after corporate tax)} = R_B \times (1 - t_c)$$

where t_c is the corporation's tax rate.