Mark Mitchell
35201: Cases in Financial Management
Section 81: Spring 2023

The Cost of Capital

> ***"Cost of capital is what could be produced by our second best idea, and our best idea has to beat it. We don't discount the future cash flows at 9% or 10%; we use the U.S. Treasury rate. We try to deal with things about which we are quite certain. You can't compensate for risk by using a high discount rate. Charlie and I don't know our cost of capital. It's taught at business schools, but we're skeptical. We just look to do the most intelligent thing we can do with the capital we have. We measure against our alternatives. I've never seen a cost of capital calculation that made sense to me."***[1]

INTRODUCTION

Firms invest in projects to generate positive future cash flows, often several years in the future. In some cases, a lengthy period may pass before these projects even begin to generate positive cash flow. To assess the feasibility of these projects, management needs to adjust the expected positive future cash flows by a discount rate, namely the project's cost of capital. Investors expect to earn risk-adjusted returns; after all, they are not investing money with the expectation of receiving sub-par returns. Estimates of the cost of capital vary over time, and across firms and industries, and also can vary at the same point in time by firms in the same industry, depending on the type of project undertaken.

The cost of capital or discount rate for a project reflects two primary components. The first reflects the time value of money; investors prefer near-term cash flows to cash flows later on. The risk-free rate reflects this time value of money. A commonly-used proxy of the risk-free rate is the expected return on

[1] Various recent quotes from Warren Buffett.

short-term government securities in well-developed economies such as the United States. The second component of the discount rate reflects the riskiness of the project. A risk premium is necessary to compensate investors for the uncertain cash flows associated with the project.

Investors do not derive the cost of capital or discount rate associated with a project in isolation; instead, investors think in terms of the opportunity cost and compare the risk of the project or firm with the risk of comparable investments. In economic terms, the return which investors expect to receive on investments of similar risk is the opportunity cost associated with the project or firm under consideration. In equilibrium, due to arbitrage, investments of equal risk will have the same expected return.

This lecture note provides a broad overview of computing the cost of capital. The first part of the note goes through the steps to calculate the cost of capital. The second part of the lecture note illustrates those steps with a brief example of the cost of capital at Northrop Grumman. Last, the Appendix provides the theoretical underpinnings to the cost of capital, as well as a broad overview of some of the basic concepts you have learned in Investments.

A caveat is in order. This lecture note assumes corporate managers behave rationally and adhere to theoretical models of asset pricing with respect to computing cost of capital estimates. In reality, most corporate managers don't necessarily adhere to models well grounded in academic theory. This is not to suggest that academia has not led to huge breakthroughs in the practice of finance, as it has on many fronts. And even if many corporate managers completely disregard practices grounded in finance theory and empirical evidence, it is still useful to understand how asset pricing models can instruct us as to the best estimate of the discount rate for project valuation.

MARKET RISK PREMIUM

There are several steps in computing the cost of capital for a project. While discount rates vary substantially across projects, a useful starting point is to assume the project has the same risk as that of the overall stock market. The required return for this hypothetical project is the expected return that the firm could make if investing in the stock market as a whole.

Eq. 1 $$\text{Discount Rate} = \text{Risk-Free Rate} + \text{Market Risk Premium}$$

Assume the project has an expected life of thirty years. Currently, thirty-year U.S. Treasury bonds are yielding 4.4% as of August 2023. Thus the discount rate should exceed at least the thirty-year Treasury yield. Investors will not hold stocks with an expected return of 4.4% over the next thirty years if they expect to receive the same return of 4.4% by investing in U.S. Treasury bonds. In other words, the market risk premium is in addition to the current risk-free rate over the expected life of the project. The true risk premium for the stock market is unknowable and thus must be estimated:

Eq. 2 $$\text{Market Risk Premium} = E(R_{MARKET}) - R_{FREE}$$

There is no universally accepted estimate for the market risk premium. Estimates range all over the place, whether in academic articles, corporate finance textbooks, investment textbooks, spreadsheet models prepared by investment bankers, finance departments in corporations, litigation consultants, judges, and many more. A commonly-used approach is to look back at how the stock market has performed relative to risk-free securities in the past. The question, however, is how far back one should go to examine the historical record. Simply using last year's stock market return relative to the risk-free rate is not sufficient. In any given year, the stock market return could be much higher or lower than the market's average historical return. And if last year's market return were less than the risk-free rate, using that rate of return in the current year to estimate the market risk premium would imply the risk is negative, which is illogical. Likewise, using just a few years of data is not sufficient either as it will lead to a very imprecise estimate of the premium. Academics and practitioners generally propose using a long historical record to obtain a reasonable estimate of the risk premium. That said, there is considerable disagreement about the length of the estimation period.

Table 1 displays the average annual rate of return and standard deviation for U.S. Treasury bills, Treasury bonds, and stocks over the 120 year period from 1900-2019.

Table 1
Average Annual Rate of Return and Standard Deviation (1900-2019)

	Return	Excess Return	Standard Deviation
Treasury bills	3.8%		2.9%
Treasury bonds	5.4%	1.6%	9.1%
Stocks	11.5%	7.7%	19.9%

U.S. Treasury bills generated an average annual rate of return of 3.8%, whereas Treasury bonds returned a higher rate of 5.4%. Stocks returned 11.5% annually, and thus have an excess return of 7.7% relative to Treasury bills, and an excess return of 6.1% relative to Treasury bonds. Assuming investor expectations turn out to be correct, we can use the realized historical excess return as a proxy for the estimated risk premium going forward. If we assume a long-term project has the same risk as the stock market, the project's cost of capital is:

Eq. 1a $$10.5\% = 4.4\% + 6.1\%$$

That is, the discount rate for the project is not the average historical return of 11.5% for the stock market; instead it is the current Treasury bond yield of 4.4%, plus the proxy for the estimated forward-looking risk premium as measured by the historical excess return of the stock market --, 6.1% --relative to bonds.

Is 6.1% a good proxy of the market risk premium relative to bonds? Conceptually, if risk aversion doesn't change through time, we should expect to see that the *ex-post* return is equal to the estimated required return, assuming we have a lengthy enough time horizon. We recognize that using too short a period of time by which to generate a proxy for the risk premium is likely to be unproductive, and less accurate, as there are often lengthy periods in which stocks underperform bonds. But it is also conceivable that risk aversion has changed over time, and that today investors are less risk-averse than a century ago when the U.S. stock market was still developing.

Influential theoretical work suggests the realized excess return is excessively large relative to predictions from economic models. If the risk premium were 6.1% during the past century, it would imply an extremely high level of risk aversion by investors, at least based on theoretical models. People invest to accumulate wealth and consume more, as opposed to investing for the sake of investing. However, consumption has low elasticity relative to the stock market. That is, when the stock market declines substantially, we do not see a corresponding decline in consumption, as would be the case if investors were extremely risk averse. The stability of consumption over time, relative to the excess return to stocks, is consistent with the notion that the realized excess return is an upward biased estimate of the risk premium.[2]

While there is valid logic in using after the fact realizations to determine estimated risk *ex-ante*, the logic assumes the draws from distributions have been as expected. But if the economy has realized far more positive draws than expected (e.g., survived multiple wars, benefitted from unexpected increases in technology, low levels of natural disasters relative to expectations, etc), then the realized returns are biased positively. Had the U.S. instead been subject to a high preponderance of negative economic impacts, stockholders could have been wiped out, or seen their investments drop precipitously, as the case in other countries.[3] The U.S. stock market is in the top half of performance by country as measured over long periods of time. Thus, the realized excess return is likely biased high. It is also the case empirically that the excess return is lower in the last fifty years relative to earlier periods.

Corporate finance textbooks tend to suggest a range of 5% to 6% in estimating the risk premium of stocks relative to long-term Treasury bonds. I assume a haircut of 0.50% to the historical excess return of 6.1%, yielding an equity risk premium of 5.6% of stocks relative to Treasury bonds. This assumed risk premium of 5.6% is right around the midpoint of what is used in practice today.[4] The haircut serves as an adjustment for the survivorship bias and for the likely reduction in risk aversion over the last 50 years versus the prior decades.

The P/E Ratio today, based on 2024 expected earnings for the S&P 500, is roughly 19.4, and Treasury bonds are currently yielding 4.4%. Thus, we can crudely calibrate the risk premium via the growing perpetuity model:

---

[2] Rajnish Mehra and Edward Prescott published a famous paper "The Equity Premium: A Puzzle" in 1985 (*Journal of Monetary Economics*) in which they were unable to reconcile the excess return to stocks with measures of consumption behavior and economic models. They initiated this research paper while visiting the University of Chicago, Mehra visiting the Graduate School of Business (Booth) and Prescott the Economics Department. Additionally, despite the enormous impact this paper has had, it took them six years to publish it. Prescott won the Nobel Prize in Economics in 2004.

[3] This is an important point which can also be made about hedge funds. When hedge funds have a particularly nasty month or quarter, they often go out of business and thus do not report their last returns to various hedge fund indexes, thereby biasing upwards the reported returns by hedge funds.

[4] I am not strongly wed to the haircut of 0.50%, rather picked it due to rounding and also to the fact that the assumed risk premium is approximate to what is used today in practice.

Eq. 3
$$V = \frac{E}{[r - g]}$$

With Treasury bonds at 4.4%, the required return is 10.0% (4.4% + 5.6%). Solving for *g* yields 3.9%. If we assume a discount rate of 10.0%, the model above implies the equity earnings of S&P 500 firms will grow 4.8% annually in perpetuity, which doesn't seem unreasonable. Suppose instead that we assume a much higher equity risk premium, such as 8.0% and thus a cost of capital of 12.4%. This would imply high perpetual growth (7.2%) of equity earnings of existing companies, which seems unlikely, given competitive pressures. And if you assume a low market-risk premium, such as 3.5% and thus a cost of capital of 7.9%, it implies equity earnings will grow at a 1.7% rate, which is less than that expected from inflation of 2.4%, as implied by the difference between Treasury bonds (4.4%) and Treasury inflation-protected bonds (2.0%).

We have little chance of knowing the true value of the market risk premium *ex-ante-*.[5] Nonetheless, having a better understanding of what the market risk premium represents, and the various complications arising in estimating it, should improve our financial decisions, as well as our decisions to invest in real assets.[6]

THEORETICAL FOUNDATIONS OF THE COST OF CAPITAL

The next step is to modify the estimate of the market risk premium to compute the cost of capital for an actual project or a firm. Generically, the cost of capital is the required rate of return for an investment project of a given level of risk. Funds used for investment projects have an opportunity cost that depends on the riskiness of the particular investment. This section briefly discusses the theoretical foundations of the cost of capital, which we will compute for firms and projects.

Investors are concerned with the tradeoff between risk and return of their portfolio. The measure of risk typically used in practice is the standard deviation. As we saw earlier, the stock market is much more volatile than the treasury market. Because of this investors demand and expect to receive a risk premium for holding common stocks, relative to investing in government securities.

While the overall stock market exhibits higher volatility than the treasury market, individual stocks are even more volatile than the overall stock market. The reason the overall market is less volatile than individual stocks in the stock market has to do with portfolio diversification. While many stocks tend to move together, especially those that are in the same industry sector, they do not always move by the same amount. A broad portfolio of stocks spread out between various industry sectors will not exhibit

[5] Perhaps ChatGPT 23.0 can instruct us as to the true risk premium!

[6] Many practitioners have far more confidence in their estimate of the market risk premium than do academics. And they often give little thought to the reliability of this estimate, perhaps not even recognizing that it is just an estimate and not a true risk premium.

movements as extreme as that of the individual stocks making up the portfolio. In finance terms, there are two types of risks from holding individual stocks:

(1) systematic or non-diversifiable risk
(2) Idiosyncratic or diversifiable risk.

Systematic risk cannot be diversified away. It stems from the fact that specific economic phenomenon, such as a recession, affects all stocks in similar ways. On the other hand, an investor can eliminate or reduce idiosyncratic risk by holding a diversified portfolio. A diversified portfolio is merely a selection of stocks that are not perfectly correlated -- that do not move in exact unison with one another. While stocks such as Goldman Sachs, Amazon, Google, Caterpillar, and American Airlines all exhibit high volatility in isolation, the risk associated with the portfolio of these stocks is lower than the risk of the individual stocks, since they are in different sectors of the economy (finance, retail, technology, manufacturing, and aviation, respectively) and thus are not always affected in the same way or to the same degree by economic shocks.

Since investors can easily diversify away idiosyncratic risk, investors cannot expect to receive a higher expected return for holding an individual stock versus holding a portfolio of stocks. However, investors cannot diversify away systematic, or market, risk, as any given stock or group of stocks will tend to vary in conjunction with the overall stock market.

The issue for the individual stock is not how volatile it is in isolation, but rather how much risk that stock contributes to the risk of one's market portfolio? The Capital Asset Pricing Model (CAPM) is the original theoretical model which addresses this question. According to the CAPM, the risk contribution of an individual stock to the total portfolio risk is proportional to the covariance of the individual stock's return with the return on the market portfolio. Beta is the CAPM term for this systematic risk of an individual stock. The beta of a stock captures the systematic relation between the movements of the stock and the movements of the overall market. If a firm has a beta of 1.0, then that firm is said to covary, or vary together, with the stock market. Public utility firms tend to have low betas -- often less than 0.50. For a stock with a beta of 0.45, when the stock market increases (or decreases) by 1.0%, it will increase (decrease) by 0.45% on average. On the other end of the distribution, luxury retailers such as Tiffany's, have high betas, often over 1.5.

To summarize the above, according to the CAPM, the equity cost of capital for a given stock or security is equal to:

Eq. 4

$$E(R_E) = R_F + \beta_E \times \left[E(R_m) - R_F\right]$$

Thus, high beta stocks have high required returns and hence high cost of equity capital, and vice versa for low beta stocks. This, of course, does not mean that high volatility stocks, such as gold mining stocks, necessarily have high betas. While gold-mining stocks exhibit very high levels of volatility vis-à-vis the typical stock, they have low betas. For example, Newmont Mining has had a beta of roughly 0.30 during the past ten years. In other words, gold-mining stocks have high idiosyncratic risk, but do not move proportionately with the overall market. It is for this reason that many investors consider gold-mining stocks as an instrument to reduce the overall systematic risk of their portfolio. Despite their high volatility, gold-mining projects have a low discount rate or hurdle rate when conducting project valuation (which can be counterintuitive).

COMPUTING THE ASSET BETA

Before the final computation of the cost of capital for a company or a project, it is necessary to adjust the equity betas to account for the capital structure of the firm or firms used to provide the equity beta estimates. The risk of common stocks, as proxied by the beta, reflects the business risk of the future cash flows and the assets held by the firm. To the extent that the firm issues debt to finance its investment opportunities, shareholders bear financial risk, as well. The more a firm relies on debt financing, the riskier its common stock. That is, borrowing creates financial leverage, and hence higher risk for the stockholders, whose claims on the cash flows of the firm are junior to the demands of debtholders.

The impact of the financial risk on shareholders via debt financing is a higher equity beta. Using the equity beta for a debt-financed firm will overstate the cost of capital estimate. Hence, the equity beta should be unlevered. Stated differently, the asset beta for the firm is a portfolio of the equity and debt betas as indicated below:[7]

Eq. 5

$$\beta_A = \beta_D \times \left[\frac{D}{D+E}\right] + \beta_E \times \left[\frac{E}{D+E}\right]$$

The underlying business risk is spread unevenly among the security holders. Most large corporations have debt that is of relatively low risk, so that the bulk of the systematic risk falls on the stockholders. As long as the D/E ratio of the corporation is low, we can assume that the debt does not have systematic risk -- that is that the beta of the debt is close to zero. By assuming that the debt beta does not have systematic risk, the asset beta formula collapses to:

Eq.6

$$\beta_A = \frac{\beta_E}{\left[1+\frac{D}{E}\right]}$$

This process is known as unlevering the equity beta to compute the asset beta. The relevant issue is calculating the systematic risk of the overall cash flows of the corporation, not just the systematic risk of the cash flows going to the stockholders.

[7] Robert Hamada, "The Effect of the Firm's Capital Structure on the Systematic Risk of Common Stocks," *Journal of Finance* (1969). Hamada combined the Modigliani-Miller Theorem with the CAPM to separate the financial risk of a levered firm from its business risk. He spent his entire career at ChicagoBooth Professor and as Dean, now Emeritus.

THE UNLEVERED COST OF CAPITAL

The cost of capital for a firm is given by:

Eq. 7 $$E(R_A) = R_F + \beta_A \times \left[E(R_m) - R_F\right]$$

This equation is the familiar CAPM formula, but here it is employed to provide the expected return for the overall firm, rather than merely for the equity of the firm: the asset or total firm beta applies to the overall systematic risk of the firm. In computing the cost of capital for a particular project for a firm, Eq. 7 will work for a project with the same risk as that for the overall firm. But projects within firms have systematic risk which can be materially different from the average systematic risk of the firm. Moreover, there are many circumstances in which equity betas are not available for the firm undertaking the project. In those situations, management should use the asset beta from publicly-traded firms with comparable risk to the project.

ESTIMATING THE COST OF CAPITAL FOR NORTHROP GRUMMAN

This section provides a brief discussion of calculating the cost of capital for Northrop Grumman, the global aerospace and defense technology firm. Note that the estimates are as of 2019. If you have access to Bloomberg, you can easily recalculate the cost of capital to see if estimates have changed.

Suppose that in planning the 2019 budget, the CFO of Northrop Grumman asks for an estimate of Northrop's cost of capital for its upcoming capital expenditures. Assume these capital expenditures will apply across the board in Northrop's various businesses and are expected to be long-term in nature, generating future cash flows which will stretch out for several years. The systematic business risk associated with these capital expenditures is identical to the systematic business risk of Northrop Grumman as a whole In essence, the exercise here is to estimate Equation 7 for Northrop Grumman.

Since the capital expenditures are expected to generate cash flows which will occur over several years, it is appropriate to use the long-term U.S. Treasury bond rate, as opposed to the short-term U.S. Treasury bill rate. The U.S. Treasury bond rate was 2.9% when the CFO of Northrop Grumman was planning the budget.[8] Thus, we assume the risk-free rate is 2.9% in calculating Northrop Grumman's cost of capital for its long-term projects. This means that the estimate we use for the market risk premium should reflect the expected return of stocks relative to U.S. Treasury bonds, as opposed to U.S. Treasury bills.[9] In practice, CFOs tend to prefer to use the U.S. Treasury bond rate as the proxy for the risk-free rate for long-term

[8] As mentioned above, the timing of this estimation is 2019 and thus the various Treasury yields are different than today in 2023.

[9] Note that the CAPM is a one-period model. In this case, the U.S. Treasury bill is much more reflective of the "true" risk-free rate than the U.S. Treasury bond is.

projects.[10] As indicated earlier, we assume, with a large standard error, the market risk premium for stocks relative to U.S. Treasury bonds on a forward-looking basis is 5.6%.

Notice that the choice of the risk-free rate and the market risk premium does not rely on the nature of the underlying project or the company itself. But the asset beta does require an understanding of the risk, specifically the systematic risk, associated with the project and the company. Again, calculating the cost of capital for Northrop Grumman's growth plans this year, will be the same as calculating Northrop Grumman's overall cost of capital, since the proposed capital expenditure is intended to apply across all of the business units at Northrop Grumman. Thus, the starting point for estimating the asset beta is Northrop Grumman's equity beta, as displayed below from a chart published in Bloomberg.

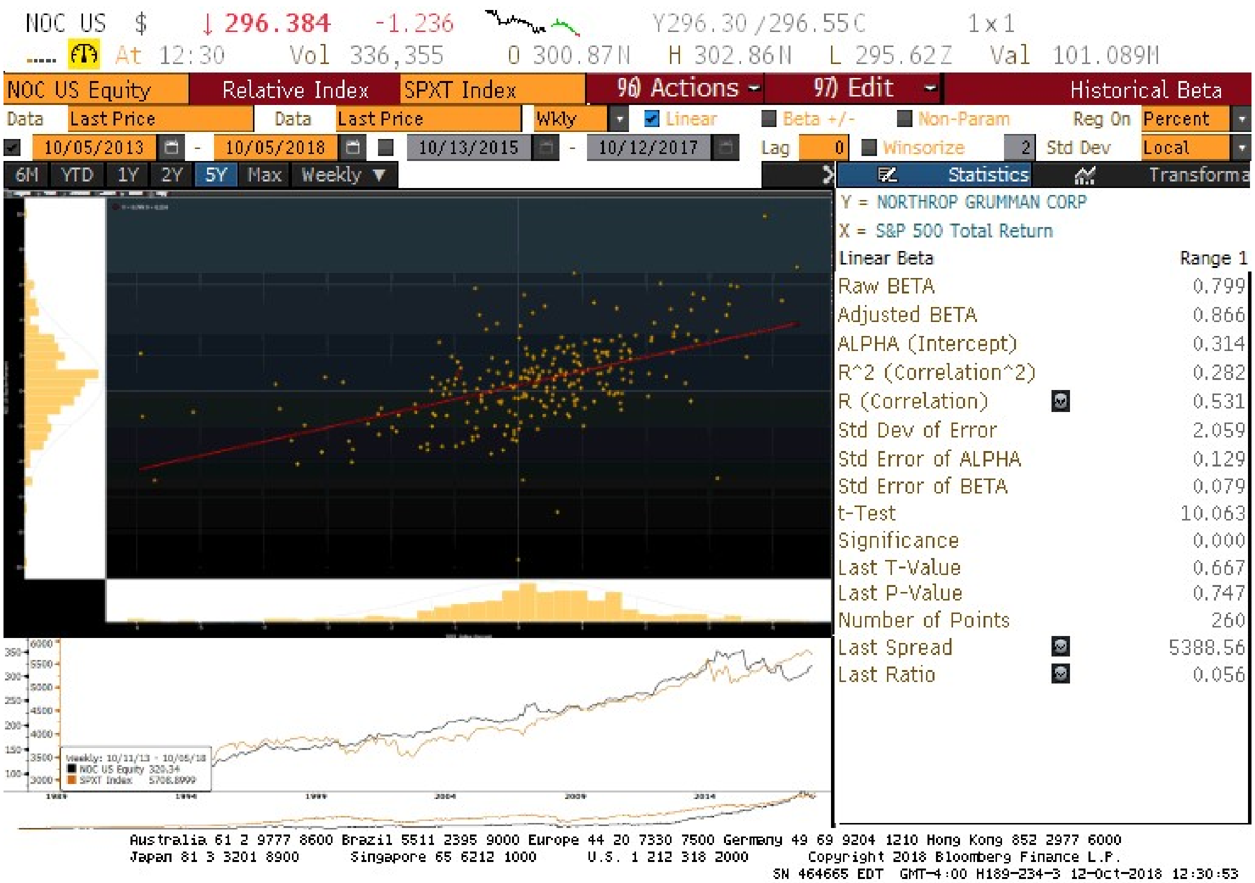

In the Bloomberg chart, the equity beta of Northrop Grumman (ticker symbol NOC) is given as 0.799 percent. To replicate the output above, the pertinent inputs are: the length of the period of time being estimated, the frequency of the return interval, the choice of market proxy, the mathematical model, etc. The period of time in this case is five years, using weekly return intervals, resulting in 260 observations. The S&P 500 Total Return (SPXT) index is the proxy for the overall stock market. Bloomberg uses linear

[10] Practitioners tend to use the U.S. Treasury security which matches the expected life of the project as the proxy for the risk-free rate. In the lecture note, *Additional Considerations in Estimating the Cost of Capital*, I discuss, at least conceptually, how to use a forecasted short-term rate for a project of any duration.

regression analysis to estimate Northrop Grumman's equity beta. With a bit of trial and error, you should be able to match the above results. Note, your return calculation for Northrop Grumman should include dividends to capture the total weekly return for Northrop Grumman, as opposed to just using returns based on Northrop Grumman's price alone.[11] Rather than computing the beta estimate for Northrop Grumman based on information from Bloomberg, you can easily download the price data (adjusted for dividends) for Northrop Grumman, and the index data for the SPXT, and calculate the beta via regression yourself if you prefer, using Excel or similar software, and obtain the same point estimate.

The objective is to generate an estimate of Northrop Grumman's forward-looking equity beta to capture its systematic equity risk. Since we do not know Northrop Grumman's forward-looking equity beta, we must use historical estimates and economic logic to calculate it. In terms of historical estimates of beta, the backward-looking period that we use as our model should have systematic risk that is similar to the future; otherwise the historical beta estimate will be too high or too low, and thus not reflective of the risk in valuing a new venture. Many factors can alter equity beta estimates over time, such as changes in the capital structure due to a share repurchase, or changes in the asset side of the business as a result of mergers and divestitures.

While this note uses weekly data over a five-year period for Northrop Grumman, it could have also used daily or monthly data, as well as periods of time of different lengths. Since Northrop Grumman's overall business hasn't changed substantially in the last few years, and because it is a liquid stock, the choice of window length and return interval is not as important as it would be for smaller stocks, or firms undergoing rapid changes. With smaller stocks, it is preferable to use weekly or monthly returns rather than daily returns, due to the somewhat stale pricing of small stocks. Alternatively, one can use daily returns and add lagged returns as well. But depending on the number of lags added, it may be more feasible to just use weekly returns.

As indicated above, linear regression analysis was used to compute the beta estimate for Northrop Grumman. But even absent the formal regression analysis, a visual inspection of the data confirms a beta estimate of a little less than one. The top chart in black displays the weekly returns to Northrop Grumman on the horizontal axis, and the weekly returns to the S&P 500 on the vertical axis. Absent the red regression line going through the data, the returns seem to line up in a reasonably straight line with a slope of less than one, with only a few large deviations. And in the white chart below it, the cumulative returns for Northrop Grumman and the S&P 500 are tracked. One measure of the reliability of the 0.799 beta estimate for Northrop Grumman is the standard error of 0.079, which is small relative to the beta estimate. That is, we can be confident that the true beta for Northrop Grumman, at least during the historical five-year period, is in the ballpark of 0.799. Of course, even absent a change in Northrop Grumman's financial structure or overall business strategy, its beta during the future could differ from its historical beta.

[11] To ensure that the Bloomberg terminal you are using is accounting for dividends and stock splits when computing returns, type PDFE (Personal Default for Equities) and hit Go. Click on the Quick Link #15 Corporate Action Settings, and make sure the boxes are checked to account for historical dividends, etc.

To the extent that Northrop Grumman has debt in its capital structure, the financial leverage will increase the equity beta, all else being equal, as discussed earlier. To estimate the asset beta, one must unlever the equity beta, and in effect, free the company of financial leverage. Using financial data from Bloomberg, Northrop Grumman's debt and equity values are displayed below over the five-year estimation period.

Table 2
Capital Structure at Northrop Grumman

	2014	2015	2016	2017	2018
Market Capitalization	29,320	34,232	40,717	53,429	52,650
Total Net Debt	2,065	4,177	4,529	4,041	13,592
Enterprise Value	31,385	38,409	45,246	57,470	66,242
Debt/Equity	6.6%	10.9%	10.0%	7.0%	20.5%

The average Debt/Equity ratio for Northrop Grumman over 2014-2018 is 11%. Assuming the debt of Northrop Grumman does not have systematic market risk, we can use Equation 6 to unlever its equity beta. [12] Given Northrop Grumman's leverage ratio is modest, and it's debt has an Investment Grade rating, we assume zero systematic market risk for the debt.

Eq. 6a

$$0.720 = \frac{0.799}{1.11}$$

To check the robustness of our equity beta ratio, one can look at industry peers. In the case of Northrop Grumman, Rockwell Collins and Harris Corporation are close competitors. Table 3 displays the equity and asset betas for all three defense contractors. You can see that Harris Corporation and Rockwell Collins have asset betas similar to the asset beta for Northrop Grumman, at 0.70 and 0.74 (compared to Northrop's 0.72).

Table 3
Equity and Asset Betas for Defense Contractors

Firm	Equity Beta	Debt/Equity Ratio	Asset Beta
Northrop Grumman	0.80	11.0%	0.72
Harris Corporation	0.96	29.6%	0.74
Rockwell Collins	0.86	22.4%	0.70

The final step is to estimate the cost of capital via Equation 7 for Northrop Grumman using the available inputs.

Eq. 7a

$$6.93\% = 2.90\% + 0.72 \times 5.6\%$$

[12] Likewise, we could have used Equation 5 to unlever Northrop Grumman's equity beta. Equation 5 is more flexible than Equation 6 as it allows for the unlevering of betas if the debt has systematic market risk. During the second half of the course when we focus on capital structure, we will more fully discuss systematic market risk of debt securities.

Thus, the discount rate which Northrop Grumman could use for long-term projects with similar systematic risk to the overall risk of Northrop Grumman is 6.93% in 2019. It reflects the opportunity cost, the expected return that an investor would receive if investing in other companies or projects of similar risk. Again, the discount rate estimate for Northrop Grumman reflects the time value of money, an overall market risk premium, and the sensitivity (beta) of the systematic project risk relative to the overall stock market.

Appendix: Theoretical Foundations of the Cost of Capital

PORTFOLIO DIVERSIFICATION

Investors are concerned with the tradeoff between the risk or variability of the returns on their portfolio, and the average or expected rate of return. The expected return on a portfolio of securities and stocks, for instance, is the weighted average of the individual expected returns on the individual securities, where the weights are the proportion of the portfolio value accounted for by each security.

Eq. A1

$$\text{Expected Return} = E\left(R_P\right) = \sum_{i=1}^{n} w_i \times E\left(R_i\right)$$

Everything else being constant, investors prefer to maximize their expected return. Investors also wish to maximize the reliability of the expected return -- to minimize the risk. In finance, we commonly use the variance or the standard deviation of the returns distribution as a measure of risk.

Eq. A2

$$\text{Variance } R_P = \sigma^2 = E[R_P - E(R_P)]^2$$

And considering a portfolio consists of a group of individual securities, the portfolio variance reflects the variance of each of the individual security returns as well as the pairwise covariances of the individual security returns:

Eq. A3

$$\text{Variance } R_p = \sigma_p^2 = \sum_{i=1}^{n} w_i^2 \sigma_i^2 + \sum_{j=1}^{n} \sum_{i=j+1}^{n-1} w_i \times w_j \times \sigma_{ij}$$

Note that as the number of securities in the portfolio increases, the first term in the variance equation above goes to zero, as *n* gets very large and the portfolio variance approaches the average covariance. That is, there are two types of risk from holding risky securities as discussed earlier, systematic or non-diversifiable risk, and idiosyncratic or diversifiable risk. Systematic risk cannot be diversified away and arises from the fact that aggregate economic phenomenon hit all stocks similarly. On the other hand, a well-diversifed portfolio can largely eliminate idiosyncratic risk. As stated earlier, a diversified portfolio is a selection of securities that are not perfectly correlated, or that do not move in perfect unison with one another.

PORTFOLIO THEORY

Using the assumption that investors are only concerned about the tradeoff between risk and expected return, portfolio theory can be used to define optimal portfolios. The first step is to identify the investment opportunity set of risky assets. The opportunity set is the set of all available investment portfolios consisting entirely of risky assets. Markowitz developed the mathematics for optimally combining the securities in the investment opportunity set into efficient mean-variance portfolios, and thus identifies the efficient frontier.[13] The term *efficient* refers to the concept of maximizing the expected return for a given level of standard deviation, or minimizing the standard deviation for a given level of expected return. For any group of stocks or securities, one can create efficient portfolios via quadradic programming to solve for the minimum variance given any expected return. Figure A1 displays the mean-standard deviation frontier for a set of securities.

Figure A1

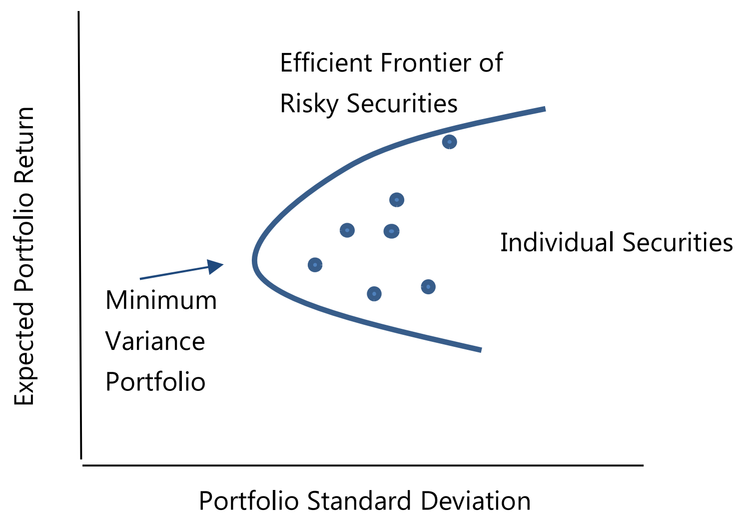

The efficient frontier is represented by the upper line sloping region of the line as it provides the maximum expected return for a given standard deviation. The points on the line dominate all of the points below the line, as well as the portion of the line in the lower region. Given that investors can choose from thousands of securities, this optimization will yield a superior, efficient frontier versus one obtained by merely considering a small sample of securities. In effect, theoretically, all investors will thus hold the same market portfolio. However, these investors will have different tolerances for risk. Some will want to own a portfolio which is less risky than the market portfolio, whereas others will want to own a

[13] Markowitz, Harry, Portfolio Selection, *Journal of Finance* (1952). Markowitz received his B.A. and Ph.D. in Economics at the University of Chicago. Portfolio Selection was his dissertation which was infamously criticized by Milton Friedman for its lack of economic insight.

portfolio which has a higher risk than the market portfolio, and thus higher potential returns. The way to achieve this is to introduce the risk-free rate to the efficient frontier as shown in Figure A2.

Figure A2

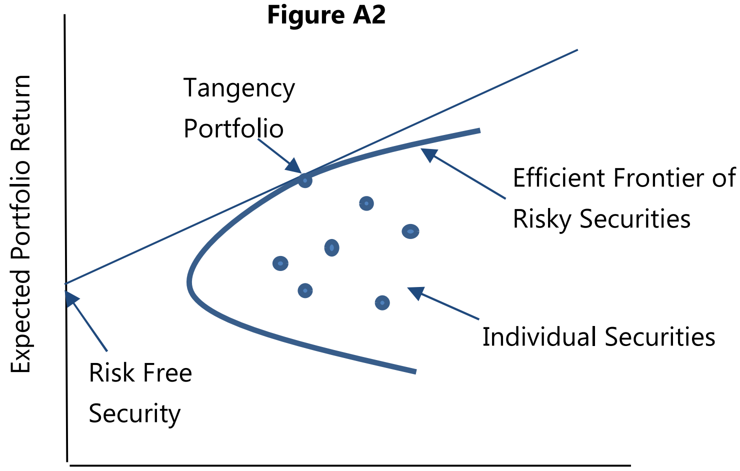

The addition of the risk-free rate improves the efficient frontier, which is now the line connecting the risk-free security and the tangency portfolio of risky assets. This line, referred to as the Capital Market Line, represents all combinations of the risk-free asset and the tangency portfolio. The tangency portfolio defines the most attractive investment opportunity set.[14] Highly risk-averse investors can invest in both the risk-free security and the tangency portfolio by investing on the line to the left of the tangency portfolio. Less risk-averse investors can borrow at the risk-free rate and lever up the tangency portfolio, thus holding a position on the Capital Market Line on the graph above to the upper right of the tangency portfolio.

The slope of the Capital Market Line from the risk-free security to the tangency point of the efficient frontier is known as the Sharpe Ratio:

Eq. A4

$$\text{Sharpe Ratio} = \frac{E\left[R_p\right] - r_f}{\sigma_p}$$

The Sharpe Ratio offers the highest risk premium per unit of risk, while the tangent portfolio maximizes the Sharpe Ratio. As noted earlier, the tangent portfolio is efficient as it lies on the efficient frontier, and when matched up with the risk-free security, it provides all combinations of the risk-free security and the tangency portfolio which are also efficient. Thus, as stated above, all investors will hold the market portfolio as it maximizes the Sharpe ratio, and based on their degree of risk aversion, will move up and down the line. Conservative risk-averse investors will hold a combination of the risk-free security and the

[14] This was Tobin's separation theorem, namely to first obtain the efficient portfolio of risky assets, and then to find the optimal fraction of risky assets and the risk-free rate to hold. James Tobin, "Liquidity Preference as Behavior Towards Risk," *Review of Economic Studies*, February 1958. Tobin won the Nobel Prize in 1981 based on his analysis of financial markets, including his separation theorem as described here.

market portfolio whereas less conservative investors may not only avoid allocating any of their portfolio to the risk-free security, instead may even borrow and lever up the market portfolio.

THE CAPITAL ASSET PRICING MODEL

Markowitz laid the foundation for asset pricing in his classic 1952 article entitled "Portfolio Selection." Bill Sharpe made an enormous additional contribution to "Portfolio Selection" in 1964 with his extension of Markowitz's and Tobin's work by incorporating economics and building a model of market equilibrium which prices securities.[15] Sharpe made several simplifying assumptions in his model, which he later coined the Capital Asset Pricing Model (CAPM):[16]

(a) all investors are price takers -- that is, there is perfect information in the market for securities
(b) it is a one-period model
(c) all investors are rational mean-variance optimizers in the spirit of Markowitz
(d) all investors have the same information and beliefs (they all use the same expected returns and covariances when estimating the mean-variance frontier)
(e) all securities are public, and investors can borrow and lend at the risk-free rate
(f) there are no transactions costs or taxes.

Sharpe's model of market equilibrium shows that all investors end up holding an identical portfolio, which is the market portfolio. That is, when investors have the same information and access to the same universe of securities, they will end up with the same efficient frontiers, and the same risk portfolio as shown in Figure A3.

---

[15] William Sharpe, "Capital Asset Prices —"A Theory of Market Equilibrium Under Conditions of Risk," *Journal of Finance* (1964). Sharpe studied under Fred Weston (who coauthored a book on mergers and acquisitions with me 20 years ago). Weston influenced him to extend the work by Markowitz. whoz acted as an unofficial dissertation advisor to Sharpe. Sharpe completed his PhD in 1961.
[16] Fama coined the term "CAPM."

Figure A3

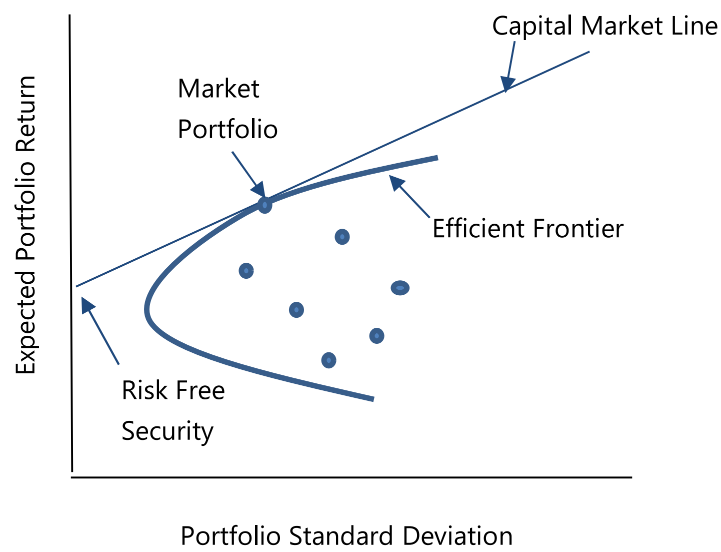

The CAPM essentially imposes market equilibrium using some specific assumptions, and identifies the tangency portfolio in the portfolio problem as the market portfolio. Each security has a weight in the market portfolio which is proportional to its market value, relative to the market value of the aggregate universe of securities. In equilibrium, individual securities are priced based on individual risk premia, which are determined by their contribution to total portfolio risk. The contribution of an individual security, *i*, to the total variance of the market portfolio is proportional to the covariance of the individual's stock return with the return on the market portfolio:

Eq. A5
$$w_i \times \mathrm{Cov}\left[R_i, R_m\right]$$

The market price of risk is the expected risk premium of the market divided by the variance of the market:

Eq. A6
$$\frac{E\left[R_m\right]-R_f}{\mathrm{Var}\left[R_m\right]}=\frac{E\left[R_m\right]-R_f}{\mathrm{Cov}\left[R_m, R_m\right]}$$

In equilibrium, if security *i* has the same risk as the market, that is, covaries perfectly with the market portfolio, then it must have the same price of risk:

Eq. A7
$$\frac{E\left[R_i\right]-R_f}{\mathrm{Cov}\left[R_i, R_m\right]}=\frac{E\left[R_m\right]-R_f}{\mathrm{Cov}\left[R_m, R_m\right]}$$

Denoting the ratio of $\mathrm{Cov}(r_i, r_m)/\mathrm{Var}(r_m)$ as β_i, Eq. A7 can be rearranged as:

Eq. A8
$$E\left[R_i\right]=R_f+\beta_i \times \left[E\left(R_m\right)-R_f\right]$$

Eq. A8 is the infamous CAPM equation. It states that the beta of a security is the right measure of risk since it is proportional to the risk that the security contributes to the market portfolio. And the risk premium of a security is the product of the security's beta and the market risk premium -- that is, $\beta_i \bullet [E(R_m) - R_f]$. Thus, the expected return to a security varies in direct proportion to its beta. The security market line (SML) portrays this relation graphically below.

Figure A4