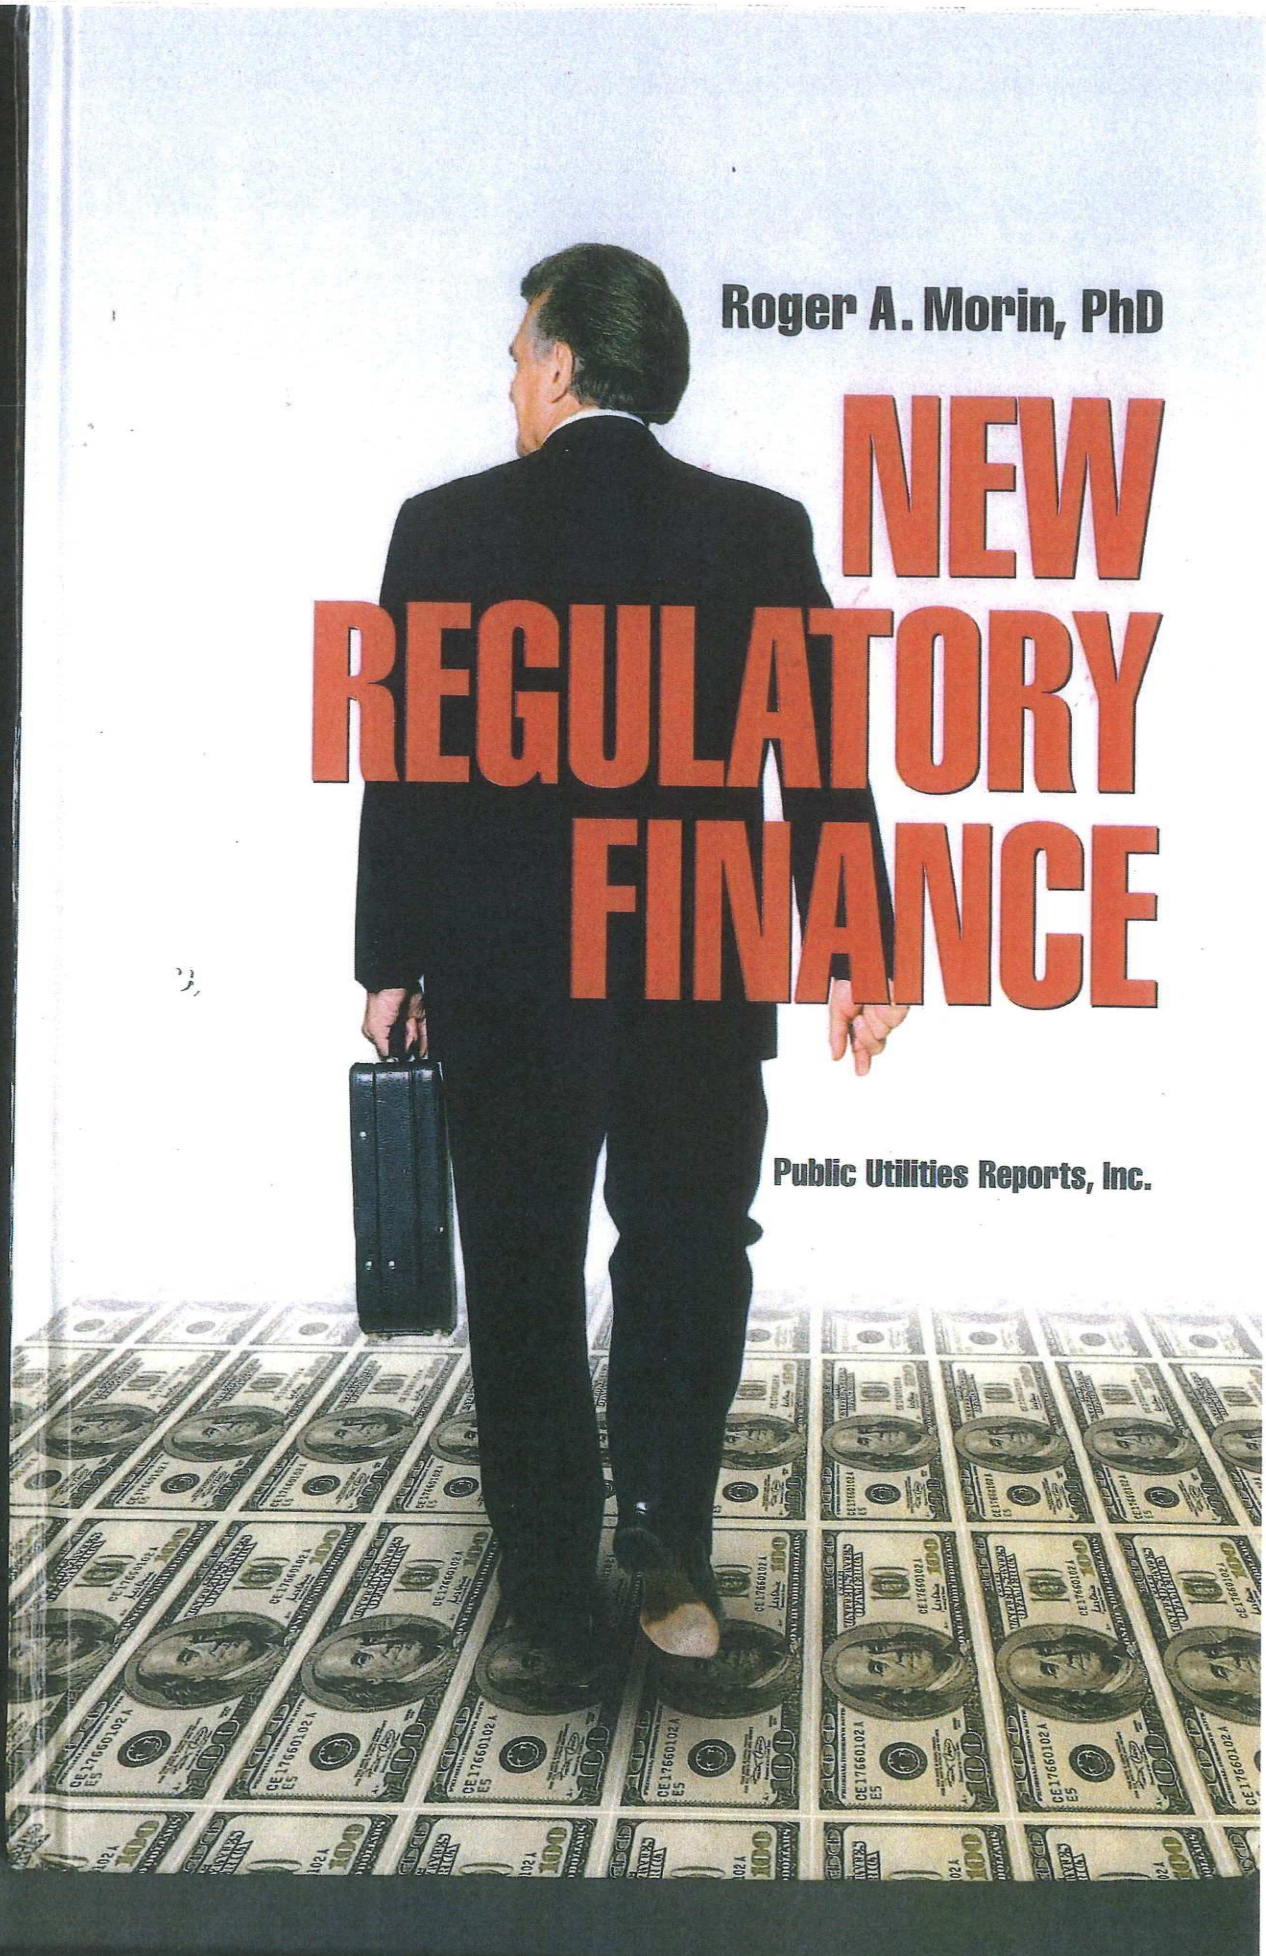
Roger A. Morin, PhD
NEW
REGULATORY
FINANCE
Public Utilities Reports, Inc.

Appendix 4-A
Arithmetic versus Geometric Means in Estimating the Cost of Capital

The use of the arithmetic mean appears counter-intuitive at first glance, because we commonly use the geometric mean return to measure the average annual achieved return over some time period. For example, the long-term performance of a portfolio is frequently assessed using the geometric mean return.

But performance appraisal is one thing, and cost of capital estimation is another matter entirely. In estimating the cost of capital, the goal is to obtain the rate of return that investors expect, that is, a target rate of return. On average, investors expect to achieve their target return. This target expected return is in effect an arithmetic average. The achieved or retrospective return is the geometric average. In statistical parlance, the arithmetic average is the unbiased measure of the expected value of repeated observations of a random variable, not the geometric mean. This appendix formally illustrates that only arithmetic averages can be used as estimates of cost of capital, and that the geometric mean is not an appropriate measure of cost of capital.

The geometric mean answers the question of what constant return you would have had to achieve in each year to have your investment growth match the return achieved by the stock market. The arithmetic mean answers the question of what growth rate is the best estimate of the future amount of money that will be produced by continually reinvesting in the stock market. It is the rate of return which, compounded over multiple periods, gives the mean of the probability distribution of ending wealth.

While the geometric mean is the best estimate of performance over a long period of time, this does not contradict the statement that the arithmetic mean compounded over the number of years that an investment is held provides the best estimate of the ending wealth value of the investment. The reason is that an investment with uncertain returns will have a higher ending wealth value than an investment which simply earns (with certainty) its compound or geometric rate of return every year. In other words, more money, or terminal wealth, is gained by the occurrence of higher than expected returns than is lost by lower than expected returns.

In capital markets, where returns are a probability distribution, the answer that takes account of uncertainty, the arithmetic mean, is the correct one for estimating discount rates and the cost of capital.

While the geometric mean is appropriate when measuring performance over a long time period, it is incorrect when estimating a risk premium to compute the cost of capital.

TABLE 4A-1
GEOMETRIC VS. ARITHMETIC RETURNS

	Stock A	Stock B
1996	50.0%	11.61%
1997	−54.7%	11.61%
1998	98.5%	11.61%
1999	42.2%	11.61%
2000	−32.3%	11.61%
2001	−39.2%	11.61%
2002	153.2%	11.61%
2003	−10.0%	11.61%
2004	38.9%	11.61%
2005	20.0%	11.61%
Standard Deviation	64.9%	0.0%
Arithmetic Mean	26.7%	11.6%
Geometric Mean	11.6%	11.6%

Theory

The geometric mean measures the magnitude of the returns, as the investor starts with one portfolio and ends with another. It does not measure the variability of the journey, as does the arithmetic mean. The geometric mean is backward looking. There is no difference in the geometric mean of two stocks or portfolios, one of which is highly volatile and the other of which is absolutely stable. The arithmetic mean, on the other hand, is forward-looking in that it does impound the volatility of the stocks.

To illustrate, Table 4A-1 shows the historical returns of two stocks, the first one is highly volatile with a standard deviation of returns of 65% while the second one has a zero standard deviation. It makes no sense intuitively that the geometric mean is the correct measure of return, one that implies that both stocks are equally risky since they have the same geometric mean. No rational investor would consider the first stock equally as risky as the second stock. Every financial model to calculate the cost of capital recognizes that investors are risk-averse and avoid risk unless they are adequately compensated for undertaking it. It is more consistent to use the mean that fully impounds risk (arithmetic mean) than the one from which risk has been removed (geometric mean). In short, the arithmetic mean recognizes the uncertainty in the stock market while the geometric mean removes the uncertainty by smoothing over annual differences.

Empirical Evidence

If both the geometric and arithmetic mean returns over the 1926–2004 data are regressed against the standard deviation of returns for the firms in the

DCF Growth Rate Check

As a reasonableness check on the DCF growth rate, the growth rate in dividends can be verified using the following relationship:[16]

Dividend Growth = Risk-free Return + Risk Premium − Dividend Yield

For example, let us say that the yield on Treasury bonds as a proxy for the risk-free return is 5%, the utility risk premium is 5.5% derived from a Capital Asset Pricing Model (CAPM) analysis discussed in earlier chapters, and the expected dividend yield for the utility industry is 4.5%. Substituting these values in the above relationship, we obtain a dividend growth expectation of 6.0% as follows:

Dividend Growth = 5.0% + 5.5% − 4.5% = 6.0%

9.6 Growth in the Non-Constant DCF Model

Although the constant growth DCF model does have a long history, analysts, practitioners, and academics have come to recognize that it is not applicable in many situations. A multiple-stage DCF model that better mirrors the pattern of future dividend growth is preferable. There is a growing consensus and ample empirical support that the best place to start is with security analysts' forecasts, that is, assume that dividend policy is relatively constant and use analyst forecasts of earnings growth as a proxy for dividend forecasts. The problem is that from the standpoint of the DCF model that extends into perpetuity, analysts' horizons are too short, typically five years. It is often unrealistic for such growth to continue into perpetuity. A transition must occur between the first stage of growth forecast by analysts for the first five years and the company's long-term sustainable growth rate. Accordingly, multiple-stage DCF models of this transition are available and were described in Chapter 8. It is useful to remember that eventually all company growth rates, especially utility services growth rates, converge to a level consistent with the growth rate of the aggregate economy.

A reasonable alternative to the constant growth DCF model is to use a multiple-stage DCF model that more appropriately captures the path of future dividend

[16] Equating the expected return from the standard DCF equation and the required return from the CAPM equation:

$$K = D_1/P + g = R_f + \text{Risk Premium}$$
$$K = D_1/P + g = R_f + \beta(R_m - R_f) \text{ from the CAPM}$$

Solving for g:

$$g = R_f + \beta(R_m - R_f) - D_1/P$$

growth than to insert a constant growth rate into the plain vanilla DCF equation. The practical challenge is to establish a reasonable growth path for future dividends. As previously discussed, an excellent starting point is security analysts' earnings growth forecasts (available from IBES, Zacks, Reuters, First Call) as a proxy for dividend forecasts. These forecasts are typically for the next five years. From the standpoint of the DCF model that extends into perpetuity, this forecasting horizon may be too short. For example, it is quite possible that a company's dividends can grow faster than the general economy for five years, but it is quite implausible for such growth to continue into perpetuity. The two-stage DCF model is based on the premise that investors expect the growth rate for the utilities to be equal to the company-specific growth rates for the next 5 years, let us say, (Stage 1 Growth), and to converge to an expected steady-state long-run rate of growth from year 6 onward (Stage 2 Growth). For example, it is quite plausible that near-term DCF growth estimates for a given company are unduly high and unsustainable over long periods, and that such growth rates are expected to decline toward a lower long-run level over time. Another example of this situation is that of companies that operate in a relatively undeveloped industry (e.g. wholesale power generation) or companies that are experiencing very high growth rates. Here again, the assumption of a constant perpetual growth rate may not be reasonable.

Blended Growth Approach

One way to account for the two stages of growth is to modify the single-stage DCF model by specifying the growth rate as a weighted average of short-term and long-term growth rates. The blended growth rate is calculated as a weighted average giving two-thirds weight to the analysts' five-year growth projections (Zacks, IBES, etc.) and one-third to historical long-term growth of the economy as a whole and/or the long-range projections of growth in Gross Domestic Product (GDP) projected for the very long term. FERC has adopted such a method in the past for determining the return on equity for gas and oil utilities.

To illustrate, two-stage DCF estimates for a group of widely traded dividend-paying diversified natural gas producers are shown on Table 9-5. Column 1 shows the spot dividend yield for each company, Column 2 shows the analyst consensus growth forecast for the next five years for each company, and column 3 shows the long-range GDP forecast of 6.5% for the U.S. economy at that time. Column 4 computes the weighed average growth, giving 2/3 weight to column 1 and 1/3 weight to column 2. Averages are shown at the bottom of the table. Adding the average blended growth rate of 9.02% to the average expected dividend yield of 2.83% shown at the bottom of Column 6 produces an estimate of equity costs of 11.85% for the group, unadjusted for flotation costs. Allowance for flotation costs to the results of Column 7 brings the return on equity estimate to 12.00%, shown in Column 7. Note realized returns over long time periods. The focus in this literature has been on the U.S. equity market because: 1) it has the most developed capital market, 2) it represents a large proportion of the international capital markets, and 3) it has long time-series of available historical data. More recently, these results have been supplemented by international analyses.

Rationale of the Historical Risk Premium Approach

Expected returns are not directly observable. As a result, realized returns are frequently used as a proxy for expected returns. This is based on the assumption that arbitrage will result in deviations between expected returns and realized returns (''surprises'') that are unpredictable and are zero-mean, that is, will cancel out, in which case realized returns provide an unbiased estimate of what returns had been expected for that period. Although realized returns for a particular time period can deviate substantially from what was expected, it is reasonable to believe that long-run average realized returns provide an unbiased estimate of what were expected returns. This is the fundamental rationale behind the historical risk premium approach. Analysts and regulators often assume that the average historical risk premium over long periods is the best proxy for the future risk premium.

Given the significant period-to-period variations in the risk premium, altering the sample period when calculating the average is dangerous because it can markedly influence the estimate. To avoid data mining, a reasonable solution is to use the entire period for which reliable data is available. Finer partitioning of the sample data, even when performed with the best intentions, raises the specter of introducing bias.

Arithmetic vs Geometric Average

One major issue relating to the use of realized returns when estimating the market risk premium from historical return data is whether to use the ordinary average (arithmetic mean) or the geometric mean return. Because valuation is forward-looking, the appropriate average is the one that most accurately approximates the expected future rate of return. The best estimate of expected returns over a given future holding period is the arithmetic average. As was thoroughly discussed in Chapter 4 and Appendix 4-A, only arithmetic means are correct for forecasting purposes and for estimating the cost of capital. There is no theoretical or empirical justification for the use of geometric mean rates of return as a measure of the appropriate discount rate in computing the cost of capital or in computing present values. In any event, the CAPM is developed on the premise of expected returns being averages and risk being measured with standard deviation. Since the latter is estimated around the arithmetic average, and not the geometric average, it is logical to stay with arithmetic averages to estimate the market risk premium. If in fact annual returns are uncorrelated over time, and the objective is to estimate the market hoped-for expected returns rather than objective required returns. Third, subjective assessments about long-term market behavior may well place undue weight on recent events and immediate prospects.

Keeping these limitations in mind, Welch surveyed finance professors on their views about the long-term equity premium in 1998 and again in 2001. The arithmetic mean long-term expected risk premium of respondents for the U.S. was 7.1% in 1998 and 5.5% in 2001. Given the deplorable behavior of equity markets in the 2000–2002 period, it would not have been surprising to see an upward reassessment of those risk premiums.

Implied Regulatory MRPs

It is instructive to examine the MRP estimates implicit in regulatory ROE decisions. The CAPM framework can be used to quantify the MRP implicit in the allowed risk premiums for regulated utilities. According to the CAPM, the risk premium is equal to beta times the market risk premium:

$$\text{Risk Premium} = \beta\,(R_M - R_F)$$
$$\text{Risk Premium} = \beta \times \text{MRP}$$

Solving for MRP, we obtain:

$$\text{MRP} = \text{Risk Premium}/\beta$$

The MRPs implied in 220 regulatory decisions for electric utilities in the United States over the period 1996–2005[12] are examined. Using the allowed average risk premium of 5.4% over that period and a beta of 0.75 for U.S. electric utilities in the above equation, the implied market risk premium is 7.2%, the same estimate as the long-term historical estimate published by Ibbotson Associates (2005). Using a beta of 0.65, the market risk premium is even higher.

Market Risk Premium: Historical or Prospective?

There are two broad approaches to estimating the risk premium: retrospective and prospective. Each has its own strengths and weaknesses, hence the need to utilize both methods.

The retrospective (historical) approach examines the historical returns actually earned from investments in stocks and bonds. Realized risk premium results are highly dependent on the choice of time period over which the security return data are compiled. Both the length of the period and the choice of end points can make a substantial difference in the final results obtained. The

[12] This study is described in more details in Morin (2005).