PRESENTATION AT CAMPUT'S

2009 ENERGY REGULATION CONFERENCE

COST OF CAPITAL

Dr. Bill Cannon

Distinguished Faculty Professor
of Finance
Queen's University School of Business

Queen's University

July 3, 2009

TABLE OF CONTENTS

	Page
1. PURPOSE OF PAPER AND PRESENTATION	1
2. THE COST OF CAPITAL CONCEPT	1
3. THE FACTORS THAT DETERMINE A FIRM'S OVERALL COST OF CAPITAL	2
4. THE FIRM'S WEIGHTED-AVERAGE, OVERALL COST OF CAPITAL	5
5. THE REGULATORY APPLICATION OF THE WACC FORMULA	7
6. THE ESTIMATION OF THE COST OF EQUITY CAPITAL	10
The Equity Risk Premium (ERP) Method	10
The Discounted Cash Flow (DCF) Method	12
The Comparable Earning (CE) Method	14
Overview of Cost-of-Equity Estimation Techniques	16
7. WHY A UTILITY'S ALLOWED RATE OF RETURN SHOULD EQUAL ITS COST OF CAPITAL	16
8. WHY A UTILITY'S ALLOWED RETURN ON COMMON EQUITY MAY BE SET SLIGHTLY ABOVE ITS COST OF EQUITY CAPITAL	18
9. CONCLUSION	20
SELECTED REFERENCES	20

COST OF CAPITAL

1. PURPOSE OF PAPER AND PRESENTATION

The purpose of this paper and my presentation at the CAMPUT 2009 Energy Regulation Conference is to convey to you an understanding of the cost of capital concept and the role it plays in the context of energy regulation. In addition, I shall review a number of the techniques and models that academics, market practitioners, and expert witnesses use to estimate the overall cost of capital and the cost of equity capital, and, in the case of the cost of equity capital, I shall talk about the pros and cons of each approach. I will also explain why, under normal circumstances, a regulated utility's allowed return should be set equal to its cost of capital, and when regulators may be justified in deviating from this rule.

2. THE COST OF CAPITAL CONCEPT

At its simplest, the cost of capital (k) is the minimum expected rate of return necessary to attract capital to an investment. The *rate of return* includes both income received during the time the investment is held plus any capital gain or loss, realized or accruing during this period, all as a percentage of the initial investment outlay. K can be viewed from both (a) a company or utility perspective and (b) from the investor's or capital provider's perspective. From the company's perspective, k is the minimum rate of return the company must promise to achieve for investors on its debt and equity securities in order to preserve their market values and, thereby, retain the allegiance of these investors. This, in turn, dictates that if corporate managers wish to protect the wealth of their shareholders and bondholders, they will invest only in projects that promise returns equal to or above their respective k's.

We are interested in the cost of capital, of course, because all utilities – private or public – at some times during their lives, must raise financial capital to pay for their investments, and both fairness and practical considerations dictate that the private and/or government investors who provide these capital funds must be adequately compensated. Raising capital is a competitive process. Private investors are under no obligation to buy a particular utility's securities, and government-owned utilities must compete with other government spending priorities. A utility will be able to secure new capital and replace maturing securities only if investors believe that they will be adequately rewarded for providing the new capital funds. That required reward, in turn, must compensate the investors for at least two things: (1) for postponing the consumption of the goods and services that they might otherwise have enjoyed had they not made the investment; and (2) for exposing their funds to the risk that they may not get all their money back or not get it back as promptly as they anticipated. The reward demanded by investors is therefore a necessary cost of doing business from the utility's point of view, just as much as the cost of labour or fuel.

For an individual investor, his or her required rate of return on an investment is a subjective, personal judgment, based on his/her tax status and degree of risk aversion, among other things. Just as beauty is in the eye of the beholder, and not some objectively measurable characteristic, so too is k at the level of the individual investor.

From the viewpoint of investors as a group, however, k can be defined more clearly and operationalized as "the expected rate of return prevailing in the capital markets on alternative investments of equivalent risk and attractiveness." There are four concepts embedded in this operational definition:

First, k is a *forward-looking* concept. Investment returns are inherently uncertain and the ex post, actual returns experienced by investors may differ from those that were expected ahead of time. K is therefore an *expected* rate of return.

Second, k is an *opportunity cost* concept. Investors have the opportunity to invest in a wide range of investments, so the expected rate of return from a given utility-company investment must be sufficient to compensate investors for the returns they might otherwise have received on foregone investments.

Third, k is adjudicated in the *capital markets*. This market price – expressed as the expected return per dollar of invested capital – serves to balance the supply of, and demand for, capital for the firm.

And, fourth, k is a function of the *risk* of the investment. It reflects the expected returns on investments in the marketplace that are exposed to equivalent risks. Another way of expressing this principle is to say that k depends on the *use* of the capital – or, more precisely, the risk associated with the use of the funds – and not on the *source* of the funds.

With respect to this fourth point, it is true that factors other than risk may influence a company's k; however, the relative risk of the company and its investments is the primary factor that differentiates one firm's k from another.

Now because k depends on the use of the funds, it is important to distinguish between the k for a company or utility, as a whole, and the k for one of its capital expenditure projects. A company or utility can be thought of as a composite of all its previous investments along with its existing investment opportunities. The k for the company will then approximate the weighted average of the k's for all its constituent investments and opportunities. Within the company, however, some of its investments may be high risk and some may be low risk. These investments will, respectively, have high and low k's, individually, and should be evaluated accordingly. That is, high-risk projects – to be worthwhile – should be expected to generate higher rates of return, while low-risk investments may be judged to be satisfactory even if they promise or produce lower returns.

3. THE FACTORS THAT DETERMINE A FIRM'S OVERALL COST OF CAPITAL

It is important to distinguish between the *determinants* of a firm's k and the *estimation* of this figure. Confusion between the two ideas arises because the k for a firm is often estimated using an approach that focuses on its weighted-average, overall cost of capital (WACC), where the "costs" associated with its various *sources of capital* – debt, preferred shares, and common equity – are averaged, using its capital structure proportions as the weights. This is the approach almost universally adopted for energy utilities during the regulatory rate-setting process. Consequently, observers are often left with the impression that the firm's overall k is *determined* by its financing proportions or sources of funds. But this is not correct. What is correct is that we may *estimate* what a firm's overall k **is**, under some circumstances, by looking at its sources of financing, but the actual *determinants* of the firm's overall k are primarily factors associated with the use of the funds it raises and, in particular, the risk attached to the use of the firm's investment capital.

So what are the factors that determine a firm's or a utility's k? These can be divided into two categories – namely, those factors that play a role in determining the k's for all companies, and those that determine the *relative* k's across companies and utilities.

The factors that play a role in determining the k's for all companies are primarily inflation expectations, the general level of interest rates, corporate income and capital tax rates, and the overall riskiness of the corporate sector. For rational investors who have the opportunity to invest their savings in real consumption and capital goods – such as houses, cottages, boats, recreational vehicles, etc. – whose prices are rising with inflation, the expected return from investing in corporate stocks and bonds must necessarily exceed the expected inflation rate, which is to say it must exceed the percentage price gains that investors could otherwise achieve by investing their savings in real goods or assets. For example, if a person invests in stocks and bonds with money that he or she would otherwise have used to buy a new car, and then the price of the car rises by 10% over the ensuing year, the person will have suffered a net financial loss unless his or her investments achieve at least a 10% return over this period.

Investors also have the option of investing their savings in default-free government treasury bills and longer-term government bonds and, if they are rational, they will not consider buying risky corporate debt and equity securities unless these latter investments promise returns exceeding government bond yields.

Furthermore, for stock and bond investments that are held outside their tax-sheltered savings plans, investors recognize that they will have to pay taxes on the income they receive from these investment holdings, and they adjust their *before-tax*, expected-return requirements upward in order to preserve their after-tax return expectations. Managers of larger corporations also recognize that any investment that they undertake that appears on their firm's balance sheet will attract federal and provincial capital taxes. Consequently, k's for these investments must be ratcheted upward to reflect these capital tax liabilities.

Finally, investors' perceptions of overall corporate-sector riskiness – and hence their return requirements for investing in corporations – can be affected by changes in their expectations about national economic growth, by exchange rate changes, and by shifts in the political, legislative, and regulatory landscape, among other things.

Therefore, any general changes in inflation expectations, interest rates, corporate tax rates, and/or corporate-sector riskiness, can be expected to change the competitive attractiveness, or *opportunity cost*, of corporate securities and thus change the average level of corporate k's.

The primary factor determining the *relative* k's across firms and utilities – and hence their deviations from the corporate-sector-average k – is the relative business riskiness of these organizations as perceived by investors. Investors are generally risk averse – they require higher expected returns to be willing to invest their money in riskier firms. The relative business riskiness of a firm, in turn, is a function of many things, including: the nature of the business it is in and the products it produces and sells; its geographic location; its cyclical sales volatility; the degree to which it uses *operating leverage* (i.e., the use of production and distribution methods generating fixed as opposed to variable costs); its level of profitability; the bankruptcy riskiness inherent in its assets and investments; the size of the firm; and the quality of its management and workers. It is worth noting that most of these factors are in some way related to the firm's assets or investments, broadly speaking, consistent with the previously-articulated notion that corporate k's are determined by the use of the funds raised, and not by the source(s) of the financing.

Beyond relative business risk, a number of other factors may affect a company's k under certain circumstances. For example, if the tradable volume of a firm's securities is small, the lack of liquidity in its securities from the perspective of large institutional investors will likely elevate its k relative to otherwise-equally-risky firms or securities. Company-specific or industry-specific tax circumstances may also, of course, favourably or unfavourably affect the attractiveness of the securities of particular companies or industries, thereby either lowering or raising, respectively, their k's. Finally, social-consciousness concerns may also raise corporate k's, as in the case of tobacco companies where institutional socially-responsible investment policies restrict investments in these firms. The effects of these non-risk factors on corporate k's tend to be isolated and relatively small, and they are not built into any of the common quantitative models used to estimate k's.

While the earlier discussion makes it clear that a company's choice of investments, or more generally its asset management policies, can be expected to have a major impact on its k, it would generally be a mistake to think that its financing or liability management policies would have a similar effect on its overall k. For the purpose of this discussion, financing policies will include selecting the firm's target capital structure proportions and the longevity of its debt and equity securities, as well as its target dividend payout ratio, where the latter effectively dictates the relative proportions of a firm's common equity financing over time that will come from retained earnings versus new share issues.

The rationale for the muted or non-existent relationship between a firm's financing policies and its overall k is this. Given the business-risk and collateral-value characteristics of a firm's assets, there is, at any time, an *optimum* set of financing policies to support these assets that is, itself, determined by the riskiness of the business and its assets. The policies are *optimal* in the sense that they are the ones that will maximize the value of the firm to its investors. Therefore, both k and the firm's financing policies are *jointly determined* by the riskiness of the firm's assets, and the financing policies themselves have no independent role to play in determining the firm's overall k. (In the next section, we shall show how financing policies apportion the risk of a firm's assets and operations among its various classes of security holders and therefore help to determine the individual k's for each type of security – however, this realization does not negate the proposition that *optimally-chosen* financing policies exert no impact on a firm's overall k.)

The only way in which a firm's financing policies can influence the firm's overall k is when the chosen or targeted policies deviate from those that are optimal from a value-maximization perspective. In this case, investors will perceive that the firm's securities are less attractive and the firm's k will rise because the firm's management has been revealed to be either incompetent or serving the interests of other stakeholders ahead of those of its investors, or both. In most circumstances where k calculations are called for, however – such as in the regulatory rate-setting arena – this scenario is not pertinent, as management is assumed to be competent and, subject to whatever constraints may exist, acting in the best interests of its bondholders and shareholders.

4. THE FIRM'S WEIGHTED-AVERAGE, OVERALL COST OF CAPITAL (WACC)

From our discussion to this point, we know that a firm's overall k reflects the required return on the firm's assets as a whole. At the same time, we know that these assets are financed by a mix of debt and equity securities and therefore this overall k must be comprised of the returns required to compensate both the firm's debtholders and its shareholders – i.e., a mixture of its cost of debt capital and its cost of equity capital. Analysts have used this realization as the conceptual foundation for developing a calculation procedure -- known as the *weighted-average cost of capital*, or WACC – for estimating a firm's overall k.

Using the WACC formula (to be detailed shortly) is an *indirect* approach to estimating a firm's overall k, as it relies on an assessment of the market-determined cost of the firm's *sources* of financing to *infer* the required rate of return on its assets, or its overall k. A *direct* approach to this task would focus on measuring the relative riskiness and attractiveness of the firm's assets in comparison with those of the typical corporation. It is the widespread use of the WACC technique that leads some people to assume, incorrectly, that a firm's k is determined by its sources of capital. While the popularity of the indirect WACC approach stems from the relative ease of obtaining its required market-based input values, it is vitally important for users of this approach to recognize that it is a *valid* overall k estimation procedure *only if* one can assume that the subject firm is targeting the value-maximizing capital structure proportions for financing its assets. Fortunately, if this lone assumption can reasonably be held to be true for the particular application – regulatory or otherwise – then the WACC approach to estimating a firm/utility's overall k turns out to be quite robust, in the sense that the approach can accommodate a wide variety of methods for estimating the input capital costs – especially the cost of equity.

The WACC formula, expressed on an after-tax basis, is as follows:

$$\text{WACC} \equiv (B/K)\, k_B\, (1 - \tau_I) + (P/K)\, k_P + (E/K)\, k_c + \tau_c$$

where: B/K = the firm's target future bond (B) or debt financing proportion, as a proportion of its total capitalization (K), expressed in market-value terms;

P/K = the firm's target future preferred share (P) financing proportion, in market-value terms;

E/K = the firm's target future common equity (E) financing proportion, in market-value terms, where E is composed of both accumulated retained earnings and the net proceeds of new equity issues;

τ_I = the firm's marginal income tax rate, expressed as a percentage;

τ_c = the firm's capital tax rate, expressed as a percentage;

k_B = the firm's average, forward-looking (not embedded), before-tax cost of bond/debt financing, expressed as a percentage;

k_P = the firm's average, forward-looking (not embedded), after-tax cost of preferred share financing, expressed as a percentage; and

k_e = the firm's forward-looking, after-tax cost of equity capital, expressed as a percentage.

Now the *theoretically correct application* of the WACC approach requires the observance of a number of important principles with respect to the choice and estimation of the input values for the WACC formula.

First, the financing weights or capital structure proportions must not only be the value-maximizing optimal ones, but they must be those that management is *targeting* for the foreseeable future, based on the firm's present and projected asset structure. Except as they may be evidence for forecasting management's intentions for the future, the firm's current, embedded capital structure proportions are irrelevant to the WACC calculation.

Second, the capital structure or financing weights in the formula must be assessed on a market-value basis and *not* on a book value or historical-cost basis. As a firm ages, the market value of its equity often deviates from the book value of the equity recorded on its financial statements. Firms perceived to have lots of profitable "growth opportunities" with modest attendant risks will often have market-determined share prices that are 2, 3, 4, or more times their accountant-recorded book values per share. On the other hand, slow-growing firms that have all-but-exhausted their chances for profitable future growth, or firms whose risk has risen dramatically (e.g., merchant energy firms post-Enron), may have shares that trade at only a fraction of their book values.

The two reasons, then, that it is imperative to use *market values* – and not book values – for the WACC calculation are: (1) investors purchase a firm's securities at their reigning market values, and not at their historical book values, and these investors gauge their required returns relative to the market values they have paid for their investment stakes, and (2) when a firm's risk, and hence its k, changes, some of this change is reflected through changes in the market values of the firm's outstanding debt and equity securities, although no changes occur immediately in the book values of these securities. Unless market values are used in the up-dated WACC calculation, this aspect of the reflection of the firm's risk change on its WACC will be overlooked.

Turning to the specification of the *capital-cost input values* appropriate for the WACC calculation, theoretical economics would call for the use of *marginal effective* costs of capital for k_B, k_P, and k_c, with the (Pareto) optimal capital structure proportions being those where the marginal effective costs of capital across all sources are equalized. Finance academics and practitioners have long realized, however, that the economist's coneeption is impractical from an implementation perspective. Instead, the finance convention is to use the *nominal* costs of debt and preferred shares in the WACC formula, where nominal corresponds to the expected returns or yields to maturity that one would read in the newspaper, based on the current market values of similar-maturity and similar-risk securities to those intended to be used by the subject firm. These nominal, future-looking costs are likely to be lower than the marginal effective cost values that economists would like to focus on. The differences between the effective marginal costs and the nominal costs are often referred to as the *hidden costs* of debt and preferred share financing.

The cost of equity capital (k_c) in the WACC formula is, on the other hand, the marginal effective cost of using incremental equity financing that reflects the degree to which the uncertainty and volatility of equity returns have been *leveraged up* by the firm's use of debt and preferred share financing. In other words, properly calculated, the k_c figure will include not only the marginal cost of equity financing, alone, but all the *hidden costs* of the firm's use of debt and preferred share financing. Another way to express this is to say that a firm's k_c is determined by both its business/asset risk exposure *and* its financial leverage risk. This was clearly illustrated during the recent worldwide financial crisis, when the k_c's of many commercial and investment banks skyrocketed – causing their share prices to plunge – largely because investors came to realize that both (a) the risks embedded in their asset holdings were far greater than originally perceived and (b) the degree of financial leverage employed by these banks was inappropriately-high for the risks that they were exposed to on the asset sides of their balance sheets.

Finally, the k_B, k_P, and k_c values should be estimates of the average *future* cost of financing from each of these capital sources and not in any way contain the *embedded* costs associated with earlier financings that are recorded on the liability side of the firm's balance sheet. These embedded costs are "water under the bridge" and have no direct relevance for the future financing costs that will have to be recovered from the earnings/returns on future investments.

5. THE REGULATORY APPLICATION OF THE WACC FORMULA

As those who have been involved in the cost of capital aspects of utility rate hearings will by now realize, there are substantial differences between the conceptually-correct, WACC calculation procedure and the approach typically followed in the hearings and reflected in board decisions. These differences – which I shall explain – are largely driven by the desire on the part of regulatory boards not to confiscate value from those investors who have provided capital to the utility in the past, when contemporaneous capital costs may have been different from those extant at the time of the hearing. While these differences flow appropriately from applying the *fairness principle* of rate regulation, they guarantee that the return on rate base derived in the typical rate hearing is only approximately related to the subject utility's true, forward-looking, overall cost of capital.

The **first** difference between the regulatory-derived WACC and the true WACC appears in the capital structure proportions used in the formula. In theory, these proportions should be those that are optimal from an enterprise-value-maximization perspective, in the light of the utility's evolving investment riskiness and the current taxation and capital market environment. In practice, of course, regulatory boards often use *deemed* capital structure proportions in their determination of a utility's WACC.

While the deemed and optimal values are likely to be similar in most cases, the forces that influence a board's choice of the deemed ratios may result in values that deviate from the theoretically-optimal ones. For example, a board's desire to keep a utility's cost of service and user rates as low as reasonably possible may nudge it toward deeming debt ratios that are unoptimally high and, correspondingly, equity ratios that are unoptimally low.

On the other hand, in order to preserve a utility's *financing flexibility* so that it may serve its growing customer base as promptly and efficiently as possible, regulatory boards typically keep an eye on the bond-rating implications of their decisions. This concern may nudge the deemed equity ratio up above its optimal level, as bond ratings are set to reflect the risk exposure of debtholders and not calibrated to reflect enterprise-value maximization. Therefore, while the combined market value of a utility's debt and equity securities may, under particular market conditions, be maximized when the debt proportion is such as to attract a triple-B bond rating from the rating agencies, a board may quite understandably deem a lower debt proportion than the optimal one – with a correspondingly higher equity ratio and overall cost of service – in order to help the utility earn a single-A debt rating and give it greater, more-assured access to new financing in order to carry out its expansion or upgrading plans.

In addition to the above considerations – which may drive a wedge between optimal and deemed capital structure proportions – regulatory practice typically results in *sticky* deemed capital structure ratios, as many boards have expressed a preference to reflect small changes in a utility's risk profile, from one hearing to the next, solely through adjustments to the equity-return award, while leaving the deemed capital structure ratios unchanged. For example, in its March 1997 *Draft Guidelines on a Formula-Based Return on Common Equity For Regulated Utilities*, the Ontario Energy Board (OEB) stated, in respect of the utilities it regulates, that "The Board, … believes that the capital structures should be reviewed only when there is a significant change in financial, business, or corporate fundamentals" (page 29). By way of illustration, if a utility's risk exposure has decreased – optimally warranting some small increase in its debt ratio – a board may choose to leave its deemed debt ratio unaltered (perhaps to placate the bond raters), but lower its equity-return award sufficiently to reflect the utility's overall lower risk exposure.

This one-variable-at-a-time adjustment procedure recognizes that there is an inherent trade-off relationship between debt ratios and the cost of equity capital. This trade-off was the subject of considerable testimony during the 1994 Multi-Pipeline Cost of Capital

Hearing (NEB RH-2-94), where the National Energy Board (NEB) chose to approve the same benchmark equity-return award for all pipelines participating in that hearing. The NEB RH-2-94 Decision stated (p.6) that "The Board is cognizant of the linkage between the rate of return on common equity and the pipelines' capital structure and has determined that any risk differentials between the pipelines can best be accounted for through adjustments to the common equity ratios rather than by making company-specific adjustments to the benchmark pipeline's rate of return on common equity". (Even so, late in the same Decision, on pages 24 and 25, the NEB expressed the view that "The determination of business risk, in our view, must necessarily involve a high degree of judgment, and the analysis is best expressed qualitatively. Under these conditions, we do not consider it realistic to refine the implications for common equity ratios to a precision of, say, one or two percent… the Board does not find it possible or meaningful to seek to determine the required equity ratio with such a degree of precision.")

Expert witness testimony in the NEB RH-2-94 hearing was inconclusive with respect to the quantification of the k_e–capital-structure trade-off except to the extent that all witnesses agreed that, for a given utility at a given point in time, there was an *inverse* relationship between its common equity ratio (CER) and its cost of equity capital. Some witnesses, however, felt that the relationship could not be quantified reliably. Others concluded that, in the vicinity of the utility's optimal CER, a one percentage point decrease in the CER (or increase in the deemed debt ratio) would occasion an increase in k_e of between 2 and 8 basis points for NEB-regulated pipelines. This relationship can be expected to change over time with the degree of risk aversion exhibited in the capital markets and with changes in corporate tax rates. Moreover, the relationship is bound to be industry specific, and the author has found evidence that suggests that k_e is more sensitive to changes in CERs for industrial companies than it is for utilities.

The **second** major difference between the regulatory and the theoretical application of the WACC formula resides in the choice of the k_B and k_P values. In theory, these cost-of-capital input figures should be entirely forward-looking estimates. In practice, however, regulatory boards incorporate the subject utility's *embedded* debt and preferred share costs into their WACC determinations, and use forecasted values only for the debt and preferred share components of the deemed capital structure that are expected to be raised during the test period. The necessity for boards to use largely embedded debt and preferred share cost rates arises from the fact that the board-determined cost of service must recover in rates a sufficient amount of money to pay the *actual* costs required to meet the utility's outstanding (i.e., past) debt and preferred share commitments – no more and no less. If this is not done, then the utility's shareholders will invariably earn a return that deviates considerably – above or below – from that which the board intends.

For the same reason, the regulatory application of the WACC formula uses the utility's average corporate tax rate – on a cash-paid or accrual basis, depending on the regulatory jurisdiction – as opposed to the theoretically-correct marginal effective tax rate. For the board to establish the before-tax, overall utility cost of capital otherwise would result in equity holders receiving unintendly high or low rates of return. Capital tax liabilities are also subsumed in the average corporate tax rate, as a matter of regulatory convenience.

The **final** difference between the regulatory and theoretical application of the WACC formula involves the estimation of the after-tax k_e value. Generally speaking, boards seek to set allow common equity rates of return somewhat above the regulated utility's true k_e, for a variety of reasons. I shall discuss the reasons for this in Section 8 of this paper after we have covered the various theoretical and practical approaches to estimating k_e.

In most jurisdictions, the traditional application of the WACC in regulatory proceedings has been a *bottom-up* exercise – that is, the costs of each of the sources of financing, along with the various (deemed) capital structure proportions, have been estimated and determined, respectively, during the course of the rate hearing, and then just plugged into the WACC formula with appropriate tax adjustments. In its Decision RH-1-2008 for the approval of 2007 and 2008 Final Tolls for Trans Quebec & Maritimes Pipeline (TQM), issued on 19 March 2009, however, the National Energy Board (NEB) approved the use of an *after-tax weighted average cost of capital* (ATWACC) for determining TQM's net revenue requirements – which is akin to a *top-down* process for determining the capital component of a utility's cost of service. While the NEB and the Alberta Energy and Utilities Board (EUB) have previously considered and rejected this top-down approach, a number of utilities are now lining up to have their ROE formulas and deemed equity ratios reviewed using the ATWACC approach, in light of the recent TQM Decision. It appears, at first blush, that the ATWACC approach provides utilities with more "wiggle room" to adjust their actual capital structure mixes as market conditions change and market opportunities emerge. It is not yet clear, however, what the advantages of the ATWACC approach are for ratepayers. Whether TQM is just a special case or the beginning of a trend is yet to be seen.

6. THE ESTIMATION OF THE COST OF EQUITY CAPITAL

There are three approaches to estimating the cost of equity capital (k_e) for regulated utilities that have traditionally been used by Canadian regulatory boards – namely, the equity risk premium (ERP) approach, the discounted cash flow (DCF) method, and the comparable earnings (CE) approach. The first two are firmly rooted in finance theory and are designed to implement the *capital attraction standard* of regulatory rate-setting. The capital attraction standard focuses on the prospective rate of return required to attract new common equity capital to a utility in competition with other investment opportunities available to investors in the marketplace. The third approach, the CE method, is not directly rooted in finance theory but is rather a reflection of the *fairness and financial-integrity standards* of regulatory rate-setting. I shall discuss each of these approaches in turn.

6.1 The Equity Risk Premium (ERP) Method

The ERP method of estimating a utility's k_e is based on the reasonable assumption that common shares are riskier, from an investor's perspective, than debt securities. Consequently, equity investors will demand a higher prospective rate of return than the contemporaneous yield prevailing on bonds to be enticed into buying shares. This extra rate of return, or premium expected return, is called the ERP. The debt rate against which this utility-specific ERP is gauged is usually taken to be the long-term government bond yield. The formula for applying this approach is:

$$k_e = r_{LTGOVT} + ERP$$

where: k_e is the estimate of the utility's cost of equity capital for the test year; r_{LTGOVT} is the forecast of the average level of the default-risk-free, long-term Canada bond yield for the test year; and ERP is the estimate of the compensatory rate-of-return premium required to attract investors to invest in the utility's shares, over the same test year horizon. The ERP value itself is gauged in comparison with the return requirements evident in the market for equity investments exhibiting investment risks similar to those of the subject utility. For this reason, the ERP method is also referred to as the "Risk Positioning" method in some texts.

In its generic form – as I have so far presented it – the ERP method leaves unspecified the investment risks for which investors need compensation. If regulatory boards were to concern themselves only with the return requirements of large, well-diversified, institutional investors, then the only type of risk for which the ERP would have to compensate would be *systematic market risk* or, as it is often called, *beta risk*. The reason for this narrowed focus is that investors, such as pension funds and mutual funds, are expected to be able to diversify away all other kinds of investment risk.

In this special situation, the ERP formula becomes what is known as the *Capital Asset Pricing Model* (CAPM), and the ERP term in the above k_e formula becomes $\beta \cdot$ MRP, where β (or *beta*) is the forward-looking estimate of the systematic riskiness of the subject utility's shares relative to the shares of the typical, publicly-traded company, and MRP, or the *market risk premium*, is the forward-looking estimate of the expected return above government bond yields that investors require to invest in the typical equity or in the stock market in general. Beta values for regulated utilities are generally thought to range between 0.40 and 0.70, while the level of the prospective MRP in North America was typically estimated to be somewhere between 3% and 5% prior to the financial market crisis that began in the summer of 2007. Since then, the MRP for North American markets has likely risen by a couple of percentage points, although it may have started to subside during the second quarter of 2009. The values for both utility betas and the overall MRP are, however, subject to considerable debate among expert witnesses.

The CAPM is only a special case of the general ERP method. For those investors in the subject utility who do not hold widely and efficiently diversified share portfolios, beta risk is an *in*adequate measure of their risk exposure, and broader concepts of risk must be subsumed in the estimate of the appropriate ERP for them. The overall uncertainty of the investment return that they may or may not receive from their ownership of the utility's shares must typically be considered and compared with the corresponding risk exposure, and excess return requirements, for the shares of other publicly-traded companies. Various numerical risk proxies and historical estimation procedures are employed for this task.

Advantages of the ERP Method

The primary advantage of the ERP method for estimating a utility's k_e is that the method focuses squarely on relative risks, which, as discussed earlier, is the principal determinant of relative costs of capital. In its typical application, the analyst looks at the risk of equities in general as compared to the risk of government bonds, as well as the risk of the subject utility's shares relative to the risk associated with other share investments.

Moreover, the ERP method is the only approach where the equity-risk impact of varying capital structure proportions can explicitly be accounted for. (For example, there are theories and formulas that relate beta values to debt and equity ratios.)

A second appealing feature of the ERP method is the reasonable expectation that individual-utility ERP values should remain fairly stable from hearing to hearing, even if there are wide fluctuations in interest rates during the intervening periods. The degree to which ERP values themselves are influenced by fluctuations in the general level of rates is a matter of some debate, however, and this factor clouds the adjustment of ERP values from one hearing to the next.

The third attraction of the ERP approach is that the data for forecasting test-year bond yields is widely available and publicly known. Indeed, forecasts of the relevant forward-looking bond yields are published by a range of investment dealers, other market participants, and various economic forecasting services, and a consensus collection of such forecasts is usually available.

Disadvantages of the ERP Method

Finding and correctly interpreting the historical equity capital costs that are often used to *infer* historical ERP values – for the market, for comparable industrials or utilities, or for the subject utility itself – are highly contentious undertakings and frequently embody either *circular reasoning* or the de facto substitution of some other k_e estimation procedure (such as the DCF method) for the ERP method. Even if the historical-average risk premium calculations are not invalidated by these procedural missteps, they tend to be highly sensitive to the estimation period chosen by the analyst and subject to considerable volatility from period to period.

In addition, there is disagreement among analysts with respect to the relevance of difference risk concepts for gauging relative ERP values, and sometimes there are disputes about the appropriateness of different risk proxies for the historical quantification of these various risk concepts – especially when the subject utility's shares are not publicly traded and historical beta risk and standard-deviation-of-investment-return values cannot be calculated directly.

Finally, as the ERP method requires a forward-looking estimate of the MRP and/or the utility-specific ERP, there is considerable debate and uncertainty about how, and indeed whether, historical evidence with respect to those variables can be used to forecast their likely future values, especially as there is considerable evidence that the prospective overall MRP is changing over time.

6.2 The Discounted Cash Flow (DCF) Method

The DCF approach to estimating a company's k_e is derived from a standard approach used to estimate a stock's intrinsic value – namely, the *constant dividend growth valuation model* developed by J.B. Williams and Myron Gordon over a half century ago. This model projects a firm's future dividend stream and then postulates that the firm's current share price in the marketplace will be the present value of these dividend payments into the infinite future, as discounted to the present at the firm's current k_c. Changes in the firm's share price can then take place only if investors' assessment of the size and growth trajectory of the firm's future dividends changes, or their k_c return requirements for the firm change, or both. By re-arranging terms in this traditional valuation model, we find the DCF-based k_c formula – namely,

$$k_c \quad = \quad E(DPS) / P \ + \ g$$

where: E(DPS) is the expected dividend payment per share over the next 12 months; P is the current market price of the firm's common shares; and g is the single percentage value that most closely approximates what investors expect the average growth rate in the firm's earning and dividend paying power to be for the infinite future.

Advantages of the DCF Method

The principal advantages of the DCF method are that it reflects investors' consensus expectations – as embodied in the current share price – and that two of the three input values are readily available in published data or may reliably be forecast from easily-accessible data. In addition, the assumptions necessary to validate the underlying valuation model seem to be reasonable ones for publicly-traded utility companies. These assumptions or requirements for validation include: payment of a cash dividend; a fairly steady growth in dividends over time; a projected growth rate in DPS payments that is less than the utility's k_c; and a valuation based on the utility's future earning power rather than, say, its asset-replacement value or liquidation value.

Disadvantages of the DCF Method

A major disadvantage associated with the practical application of the DCF method is the uncertainty surrounding, or speculative nature of, the forecast of the utility's infinite future growth rate. First of all, a regulatory board's equity-return awards over time will impact a utility's experienced growth rate which, if used in the estimate of the growth rate factor in the DCF formula, will create *circularity* in the board's reasoning process by making its future k_c determinations a self-fulfilling prophecy of its previous equity-return awards. Second, a utility's historical growth experience may not necessarily be indicative of its actual or expected future growth rates. Finally, forecasting a single-number growth rate value to infinity is a highly problematic task when a utility may go through phases of both rapid and slow growth in future years, before maturing and possibly ceasing to grow altogether. Some analysts try to address this last weakness by using what is called the "Two-Stage Gordon Model", where the future is divided into two segments – an initial time frame where dividend growth may reasonably be forecasted based on presently-known conditions, and the remaining time to infinity after the initial period. While this allows for a more-normal growth pattern to be envisioned, it still requires the analyst to make dividend growth rate projections – two of them – that may or may not be realistic reflections of investors' expectations.

Another limitation of the DCF method arises when the shares of the subject utility are not publicly traded and, hence, there is no direct evidence of its investor-determined share price or of its stand-alone dividend policies. Much of the appeal of the DCF approach is then lost if dividend yield values for the formula have to be inferred from "comparable" utilities or from non-regulated-firm data, and subject-utility-specific risk changes cannot be factored into the k_e estimation. Complicating this is the fact that the DCF model does not incorporate any relative risk factor directly into its model structure, but merely assumes that risk-related differences among firms will be adequately reflected in share prices.

Some of the assumptions or requirements of the underlying DCF valuation model also seem at odds with both reality and the regulatory environment. For example, the underlying model is premised on the firm's having a constant k_e over time, while the regulatory process envisions that a utility's k_e and allowed equity return may change over time. In addition, the underlying model assumes a constant dividend payout ratio over time, which may not be appropriate for a utility which is transitioning from one growth phase to another. Finally, the model requires that stock prices be in equilibrium relative to investor long-run risk assessments and growth rate expectations. But this assumption is fallacious if investor assessments are in transition themselves, or stock prices reflect short-run guesses with respect to, say, up-coming regulatory decisions.

6.3 The Comparable Earning (CE) Method

The CE method is based on the *fairness standard* of regulatory rate-setting whereby it may be reasonable for an investor to expect to receive a return on the book value of his/her investment in the shares of a utility that is equal to that being received by investors in other regulated and non-regulated companies of similar risk and operating characteristics – that is, equal to the achieved returns on *comparable companies*. The legal basis for this fairness standard is reviewed in Section 7 of this paper.

The identification of a sufficiently large sample of comparable companies may comprehend comparisons with the subject utility along such dimensions as industrial classification, size, corporate longevity, return volatility, dividend policies, credit ratings, and the presence of monopoly effects on achieved returns. An appropriate time period is chosen for calculating historical average returns on equity (ROEs). Once historical average ROEs for the sample of comparable firms are determined, adjustments may be required (a) to account for any differences in risk between the subject utility and the typical firm in the sample of comparable firms and (b) possibly to account for how test-period economic conditions may differ from those during the historical-ROE-averaging period.

Analysts often apply the *financial integrity standard* of regulatory rate-setting in conjunction with the CE method. (The legal basis for the financial integrity standard is set out in Section 7.) The financial integrity standard is usually interpreted to mean that return awards should not be set so low that, were the utility required to issue new common equity, the new issue would cause the book values of the shares held by the utility's pre-existing shareholders to be diluted – effectively confiscating some of the value originally contributed by these investors. The implication of this is that looking at the experienced ROEs for comparable firms alone often tells us very little about the adequacy of these returns in shareholders' eyes. If allowed and actual returns are more generous than is necessary to satisfy the financial integrity standard, then investors will bid up the price of the firm's shares and these shares will normally trade at a market-value-to-book-value ratio (MV/BV ratio) considerably greater than one. If, on the other hand, returns are consistently inadequate in the light of investors' expectations and alternative investment opportunities, investor reaction will cause the firm's share price to fall and the shares will typically trade at a MV/BV ratio less than one. In this latter case, the firm will not likely be able to issue new common shares without impairing the book value of the shares of existing shareholders. Consequently, in the context of the CE method, analysts must be cognizant of the average MV/BV ratios for the comparable firms, as well as average realized rates of return on equity, when gauging the suitability of these returns as a benchmark for setting allowed utility ROEs. A compensating adjustment to the historical average ROEs for the comparable firms may be called for if the award of these average returns would drive the subject utility's MV/BV ratio to an unacceptably high or low level.

Advantages of the CE Method

The primary advantage of the CE method is that it relies on easily accessible, actual financial data that investors commonly use to assess the comparative profitability of firms. Moreover, the averaging-of-historical-ROEs procedure is easy for regulators to understand and the approach does not rely on some theoretical financial model, whose underpinnings and applicability may be hard for regulators to judge.

Disadvantages of the CE Method

There are numerous shortcomings of the CE method as an approach to estimating the k_e for either a regulated or a non-regulated firm. The first of these is that the CE method is not even *intended* to measure a firm's current k_e. As the CE approach is based on book-value returns and not the investment returns experienced by shareholders, the CE method can only be expected to yield an estimate of a firm's k_e if **both** (a) the market value and book value of the firm's shares are always equal and (b) historically-achieved rates of return are equal to the investment rates of return investors expect or require going forward. For the vast majority of possible comparable companies, neither of these conditions is likely to hold.

Beyond this theoretical limitation, there are numerous practical problems that analysts experience when they apply the CE method. There are often difficulties in assembling an acceptable sample of comparable companies against which to assess the subject utility. The selection process is subjective and open to bias and conflicting interpretations. The same criticisms apply to the analyst's choice of the suitable time period from which to draw historical ROE and risk evidence.

Another weakness is that most of the data used in the CE method is historical in nature, and there is no assurance that the future performance of the comparable firms will mirror their historical results – and provide a useful approximation of the future opportunity cost for utility shareholders – if underlying economic conditions impacting on inflation and corporate profitability are changing. Moreover, the adjustments that analysts make to address this problem are typically subjective and lacking in theoretical support.

Finally, as has become painfully apparent over the past decade, the basic ROE data for the comparable companies themselves may be inappropriate for the purpose of the CE analysis because of the wide divergences between accounting earnings and what economists refer to as *real economic earnings* – that is, the earnings calculation that is most relevant to investors. In particular, the kinds and degrees of bias introduced by standard accounting conventions are likely to be significantly different for regulated utilities as compared with the typical, otherwise-comparable, industrial or resource company. (Note, however, that side-stepping this problem by employing other regulated firms as comparables introduces the problem of *circularity*, where one board's equity return award establishes the historical ROE evidence that influences awards in other jurisdictions.) In addition, there is widespread evidence that many corporations have systematically misrepresented their true accounting earnings in order to paint a more favourable picture for investors and thereby sustain, or "juice up", their share prices. These last two problems undermine the legitimacy of using accounting ROE values to estimate corporate k_e's even when equity book values and market values are approximately the same.

6.4 Overview of Cost-of-Equity Estimation Techniques

The three k_e estimation techniques reviewed above all have their strengths and weaknesses, and all potentially have a role to play in regulatory rate hearings. If, however, the analyst's purpose is strictly to estimate the utility's forward-looking k_e, then the ERP and DCF methods will trump the CE approach. And, as between the former two techniques, the ERP method is likely to be more useful and reliable than the DCF approach, especially if the utility's risk profile is changing over time and the shares of the stand-alone regulated entity are not publicly traded. This preference for the ERP method is the one that the NEB, the OEB, and many other Canadian energy regulators have effectively made with their adoption over the past decade a formulaic approach to setting and annually adjusting allowed equity returns for the regulated utilities and pipelines within their jurisdictions.

7. WHY A UTILITY'S ALLOWED RATE OF RETURN SHOULD EQUAL ITS COST OF CAPITAL

There are both legal and economic reasons why regulators, in carrying out their responsibility to set "just and reasonable" rates, should strive to set a regulated utility's allowed rate of return equal to its cost of capital. In Canada, the decision of the Supreme Court of Canada in *Northwestern Utilities Ltd. vs the City of Edmonton* (1929 SCR192) is often cited to support the equating of the *opportunity cost of capital* with a *fair return* to the utility and its investors. Mr. Justice Lamont, in his judgment in the Northwestern Utilities case, stated that:

> "By a fair return is meant that the company will be allowed as large a return on the capital invested in its enterprise, which will be net to the company, as it would receive if it were investing the same amount in other securities possessing an attractiveness, stability and certainty equal to that of the company's enterprise."

In the U.S., the courts have established that investors in rate-regulated enterprises must be allowed an opportunity to earn returns that are both (a) sufficient to attract new capital and preserve creditworthiness and (b) comparable to those that they would expect in unregulated companies exposed to the same degree of risk. The seminal decisions in this regard are *Bluefield Water Works Co. vs Public Service Commission*, 262 U.S. 679 (1923) and *Federal Power Commission vs Hope Natural Gas Co.*, 320 U.S. 591 (1944). The pertinent part of the opinion in the *Bluefield* case, written by Justice Butler, states:

> A public utility is entitled to such rates as will permit it to earn a return on the value of the property which it employs for the convenience of the public equal to that generally being made at the same time and in the same general part of the country on investments in other business undertakings which are attended by corresponding risks and uncertainties; but it has no constitutional right to profits such as are realized or anticipated in highly profitable enterprises or speculative ventures. The return should be reasonably sufficient to assure confidence in the financial soundness of the utility, and should be adequate, under efficient and economical management, to maintain and support its credit and enable it to raise the money necessary for the proper discharge of its public duties.

The pertinent passage from the *Hope* decision states:

> The rate-making process under the act, i.e., the fixing of "just and reasonable" rates, involves a balancing of the investor and the consumer interests…the investor interest has a legitimate concern with the financial integrity of the company whose rates are being regulated. From the investor or company point of view it is important that there be enough revenue not only for operating expenses but also for the capital costs of the business. These include service on the debt and dividends on the stock … By that standard the return to the equity owner should be commensurate with returns on investments in other enterprises having corresponding risks. That return, moreover, should be sufficient to assure confidence in the financial integrity of the enterprise, so as to maintain its credit and to attract capital.

Taken together, these court decisions provide the basis for equating the legally fair return to a utility's investors with its cost of capital. This correspondence is evident from the fact that the opportunity cost of capital for a regulated utility is, by definition, the expected return that investors could anticipate from other investments possessing an equivalent level of risk, and from recognizing that investors will not provide capital to the utility unless these investments are expected to earn their costs of capital.

From an economic perspective, there are both direct, microeconomic benefits and indirect, macroeconomic benefits from a regulator's setting the allowed rate of return for a utility equal to its cost of capital. The direct benefits are that the utility's customers will pay the lowest cost of service in the long run if the utility's investors expect the allowed rate of return to equal the cost of capital. On the one hand, if regulators set the allowed return *above* the cost of capital, then the utility's shareholders will earn more than they could elsewhere on equally-risky investments. The utility's customers will pay for this overly-generous return, but could have received the same service without paying this extra amount. In effect, there will be an unnecessary transfer of wealth from the utility's customers to its shareholders, where the latter are likely to see the price of their shares rise as result of the receipt of this excess return. On the other hand, if the allowed return is set *below* the utility's cost of capital, its shareholders will earn less than they might otherwise on comparable-risk investments. While current utility customers may gain temporarily from this situation – and experience a wealth gain at the expense of the utility's shareholders, whose share values are bound to fall – investors will balk at financing the future investments that the utility needs to replace, upgrade, and expand its rate base to maintain an adequate level of service for its customers. Customers will eventually suffer from inferior service, which is as much a *cost* to them as excessive rates. The lesson is that the regulator should set the allowed return equal to the cost of capital so that the utility can attract the necessary financing to achieve the optimal amount of rate base investment at the minimum effective cost to ratepayers.

There are also indirect or macroeconomic benefits from equating utility allowed returns with their costs of capital. In our competitive capital markets, the cost of new capital financing is a good measure of the value of that capital in alternative uses. If regulators allow utilities to earn returns higher than their costs of capital, there will be an incentive for investors to direct an unproductively-high level of investment toward the regulated sector and divert capital at the margin from sectors where investments may be more productive. Conversely, if regulated returns are too low and fail to compensate investors for the risks they face, capital that should have been employed in the regulated sector may be diverted into less productive investment. At a macroeconomic level, society as a whole suffers from the misallocation of scarce resources when regulated firms are allowed to earn more than, or are forced to accept returns lower than, their true costs of capital.

It is often said that regulation should substitute for competition in the sense that the regulator should attempt, by its actions, to incent the regulated utility to make the pricing, output, and investment decisions that would be forthcoming under competitive conditions. By setting the utility's allowed return equal to its cost of capital, customers will pay the true cost of providing the utility's service, as they would under perfect competition, and their service demands will elicit an economically efficient allocation capital between regulated and unregulated industries.

8. WHY A UTILITY'S ALLOWED RETURN ON COMMON EQUITY MAY BE SET SLIGHTLY ABOVE ITS COST OF EQUITY CAPITAL

Despite the strong economic case for equating a regulated utility's allowed return with its overall cost of capital, there are a number of *fairness* and *financial integrity* related considerations, along with some practical factors, that may legitimately persuade a regulatory board to allow a utility to earn a return on the common-equity-financed portion of its rate base somewhat above its model-determined, "bare bones" cost of equity (k_e). The "bare bones" k_e is the value that is derived when the analyst employs either of the ERP or DCF estimation techniques.

We shall look at the practical factors first. The ERP or DCF model-derived k_e estimates reflect the required returns for the utility's current shareholders, but incorporate no allowance for the *flotation costs* – such as underwriting and legal fees, and market pressure effects – that are incurred if the utility has to issue new common shares to raise capital. In the situation where a utility is expected to access new external equity periodically over time, boards have seen fit to add an allowance for flotation costs on top of k_c when setting the utility's allowed return on equity (ROE).

Preservation of the utility's *financing flexibility* is another practical consideration that may persuade a regulator to set the utility's allowed ROE slightly in excess of its "bare bones" k_c. Allowed ROEs are typically set no more frequently than once a year and should be sufficient to allow the utility to attract new equity capital at any time during the year, as long as the utility's risk profile is stable and capital market conditions remain "normal" or "accommodative" throughout the year. However, the situation may arise where it is desirable – from the board's or ratepayers' perspective – for the utility to undertake an investment which appropriately requires equity financing at a time when market conditions have deteriorated and the utility's shares are trading at a price below their book value. Then the utility managers will face the choice of either postponing the investment, perhaps to the detriment of ratepayers, or going ahead with the investment and selling new equity a below-book-value price. In the view of some analysts and regulators, issuing new shares at a below-book-value price is seen as undermining the *financial integrity* of the utility's shares, which may impair the utility's credibility in the capital markets and compromise its future access to financing. It may also be seen as *unfair* that the new shareholders will get their shares at a price lower than that paid by investors in previous years and, in doing so, diminish or dilute the book value per share of the holdings of previous equity investors. Consequently, to reduce the chance that utility managers will be forced to issue new shares at prices below book value and undermine the financial integrity of their shares, regulators may set the allowed ROE at a level modestly above the utility's k_c, with the expectation that this will cause the utility's shares to trade at prices modestly above their book value, under normal conditions, and at no less than book value under less favourable conditions – thus enhancing the utility's financing flexibility.

Concern for a utility's credit ratings is another practical consideration that may call for regulators to set the allowed ROE above the utility's k_c. While utility rate-setting and the awarding of allowed ROEs typically follow annual cycles, regulators must recognize that the creation and preservation of corporate credibility in the capital markets is a longer-term process – but still one that is vital to serving the interests of both investors and ratepayers. So, if a utility's cost of equity capital falls rapidly because, say, interest rates fall sharply, blindly setting the utility's allowed ROE equal to its current k_c may cause its interest coverage ratio and other credit metrics to fall to levels that, if left uncorrected, will jeopardize its bond ratings. In this situation, a board may very well choose to postpone the full reduction of the utility's allowed ROE – toward its new k_c – in order to keep its prospective coverage ratios at an acceptable level and protect the *financial integrity* of the utility as a whole – thus preserving its unfettered access to capital market financing in future years. What may be seen as "unfair" to this year's ratepayers will help ensure the maintenance of service levels for ratepayers in future years.

Fairness and financial integrity considerations at a more philosophical level may also occasionally call for a board to set a utility's allowed ROE at a level above its k_c. Generally, setting the allowed ROE equal to the utility's k_c satisfies the usual notions of regulatory fairness, as it is "fair" to shareholders that they should receive a return equal to that which they would otherwise have earned on comparably-risky foregone investments, and it is "fair" to ratepayers that they should be protected from having to pay for the excessive monopolistic returns that utility shareholders would otherwise be able to extract from them outside of the regulatory framework.

Nevertheless, a return that is "fair" to current and recently-attracted shareholders may be seen as "unfair" to those shareholders who invested in the utility at an earlier stage. Suppose, as an extreme example, the original shareholders in a fledgling utility facing highly uncertain prospects were enticed into purchasing the new utility's shares by the prospect of elevated, risk-compensating, allowed returns into the foreseeable future. Years later, however, when the utility has grown and stabilized and its risk profile has been lowered, regulators may reduce its allowed ROE. While the lowered ROE may very well satisfy new shareholders – who are now taking on little investment risk – it may be seen as "unfair" and confiscatory from the perspective of the original shareholders who *did* once bear considerable risk but now see their deserving reward for this burden cut short before they have been fully compensated.

9. CONCLUSION

While the cost of capital *concept* is reasonably easy to grasp, understanding its role in the regulatory process and the subtleties of its calculation in specific and changing circumstances is far from easy or straight forward. I hope that this paper and presentation have helped to illuminate the role and importance of the cost of capital concept in the regulatory environment and, if nothing else, given you an appreciation for the degree of complexity and judgment that are inevitably involved in estimating a utility's cost of capital as a precursor to establishing its allowed return on equity.

SELECTED REFERENCES

Bonbright, James, C., *Principles of Public Utility Rates*, Columbia University Press, New York, 1961.

Draft Guidelines On a Formula-Based Return on Common Equity For Regulated Utilities, Ontario Energy Board, Toronto, March 1997.

Kolbe, A.L., Read, J.A., and Hall, G.R., *The Cost of Capital: Estimating the Rate of Return for Public Utilities*, The MIT Press, Cambridge, Mass., 1984.

Morin, Roger A., *Utilities' Cost of Capital*, Public Utilities Reports, Inc., Arlington, Va., 1984.

Ross, S.A., Westerfield, R.W., Jordan, B.D., and Roberts, G.S., *Fundamentals of Corporate Finance*, Fifth Canadian Edition, McGraw-Hill Ryerson, Toronto, 2005.

1994 Multi-Pipeline Cost of Capital Hearing, "Reasons for Decision", RH-2-94, National Energy Board, March 1995.