A PUBLIC POLICY PRACTICE NOTE

Selecting Investment Return Assumptions: Considerations When Using Arithmetic and Geometric Averages

July 2019

American Academy of Actuaries
Pension Committee

American Academy of Actuaries
Objective. Independent. Effective.™

A PUBLIC POLICY PRACTICE NOTE

Selecting Investment Return Assumptions: Considerations When Using Arithmetic and Geometric Averages

July 2019

Developed by the Pension Committee
of the American Academy of Actuaries

The American Academy of Actuaries is a 19,500-member professional association whose mission is to serve the public and the U.S. actuarial profession. For more than 50 years, the Academy has assisted public policymakers on all levels by providing leadership, objective expertise, and actuarial advice on risk and financial security issues. The Academy also sets qualification, practice, and professionalism standards for actuaries in the United States.

2019 Pension Committee

Bruce Cadenhead, MAAA, FSA, EA, FCA, Chairperson
Elena Black, MAAA, FSA, EA, FCA, Vice Chairperson

Michael Antoine, MAAA, FSA, EA
Mike Bain, MAAA, ES, FCA, FSPA, ASA
Rachel Barnes, MAAA, FSA, EA, CERA
Tim Geddes, MAAA, FSA, EA, FCA
Stanley Goldfarb, MAAA, FSA, EA
Scott Hittner, MAAA, FSA, EA, FCA
Lloyd Katz, MAAA, FSA, EA, FCA
Grace Lattyak, MAAA, FSA, EA, FCA
Tonya Manning, MAAA, FSA, EA, FCA
A. Donald Morgan, MAAA, FSA, EA, FCA
Nadine Orloff, MAAA, FSA, EA, FCA
Jason Russell, MAAA, FSA, EA
James Shake, MAAA, EA, FCA
Mark Shemtob, MAAA, FSA, EA, FCA
Mary Stone, MAAA, FSA, EA, FCA
Aaron Weindling, MAAA, FSA, EA, FCA

The Committee gratefully acknowledges the contributions of David Kausch, Gerard Mingione, Mitch Serota, and former Pension Committee Chairperson Michael Pollack.

American Academy of Actuaries
Objective. Independent. Effective.™

1850 M Street NW, Suite 300
Washington, DC 20036-5805

TABLE OF CONTENTS

INTRODUCTION

This practice note is not a promulgation of the Actuarial Standards Board, is not an actuarial standard of practice (ASOP) or an interpretation of an ASOP, is not binding upon any actuary, and is not a definitive statement as to what constitutes generally accepted practice in the area under discussion. Events occurring subsequent to the publication of this practice note may make the practices described in the practice note irrelevant or obsolete.

This practice note was prepared by the Pension Committee of the Pension Practice Council of the American Academy of Actuaries to provide information to actuaries on current and emerging practices in the selection of investment return assumptions based on anticipated future experience. The intended users of this practice note are the members of actuarial organizations governed by the ASOPs promulgated by the Actuarial Standards Board.

This practice note may be helpful when setting investment return assumptions, or providing advice on setting investment return assumptions, for funding (where permitted by law) and for financial accounting in connection with funded U.S. benefit plans. It does not cover the selection and documentation of other economic assumptions or demographic assumptions.

The Pension Committee welcomes any suggested improvements for future updates of this practice note. Suggestions may be sent to the pension policy analyst of the American Academy of Actuaries at 1850 M Street NW, Suite 300, Washington, DC 20036 or by emailing pensionanalyst@actuary.org.

BACKGROUND

Actuarial Standard of Practice No. 27 (ASOP No. 27), *Selection of Economic Assumptions for Measuring Pension Obligations*, provides guidance to actuaries in selecting economic assumptions such as those relating to investment return, discount rates, and compensation increases.

Key provisions of ASOP No. 27 relating to the determination of investment return assumptions include the following:

- Assumptions should be reasonable and consistent with other economic assumptions selected by the actuary for the measurement period (Sections 3.6 and 3.12).

- Assumptions should reflect the actuary's observations of the estimates inherent in market data and/or the actuary's estimate of future experience (Section 3.6[d]).

- Assumptions should have no significant bias (Section 3.6[e]).[1]

- The actuary should review appropriate recent and long-term historical economic data as part of the assumption-setting process (Sections 3.4).

- Active management premiums should not be anticipated without relevant supporting data (Section 3.8.3[d]).

Complex issues arise in the determination of investment return assumptions, especially for an investment return assumption that will be used as a discount rate (i.e., as a means for determining the present values of promised benefit payments payable over long periods). In particular, the ASOP acknowledges the distinction between assumptions that reflect arithmetic versus geometric average returns (Section 3.8.3[j]). Arithmetic averages generally exceed geometric averages, but some issues and concerns may arise in developing investment return assumptions based on these higher rates. The ultimate choice between these approaches, or the adoption of an alternative approach, will likely depend on purpose of the measurement. The approaches may produce materially different results.

This practice note provides discussion and background information relating to this technical issue. It focuses primarily on considerations relating to the use of *the return assumption as a discount rate*; other situations are noted in an appendix. The body of the practice note is divided into seven sections:

I. Terminology: sets forth definitions of terms that will be used frequently. Readers are encouraged to review this section carefully, as usage in this practice note may differ slightly from what may be used in other contexts.

II. Example: demonstrates geometric and arithmetic computations for historical performance.

III. Forecast Models—The Effect of Uncertainty: shows how return measures are affected by the variability of outcomes.

IV. Relationships Among Statistics: compares means and medians in the context of arithmetic and geometric models.

V. Analysis of Forecast Returns: addresses stochastic simulations and analysis of results.

VI. Considerations for Actuaries: presents additional issues to be considered with respect to discount rate selection.

[1] The ASOP contains an exception "when provisions for adverse deviation or plan provisions that are difficult to measure are included and disclosed under section 3.5.1, or when alternative assumptions are used for the assessment of risk."

VII. Conclusions: summarizes the key points from the practice note.

The material presented in this practice note is complex and technical. Although an initial read-through may not require a major time investment, actuaries may find it beneficial to devote several hours to a more in-depth review and study of the concepts, arguments, and applications presented. The practice note offers three appendices and a bibliography to support further independent study.

I. TERMINOLOGY

Setting an investment return assumption can require the application of concepts that are highly technical and involve subtle theoretical distinctions. Gaining a thorough understanding of these concepts may be challenging because different authors may use terminology differently. Some terms can also be used in a less technical sense in other contexts, and they therefore might have developed certain general connotations that can be confusing or misleading when those terms are used in a technical setting.

Accordingly, this section lays out the terminology that will be used throughout this practice note. Note that this terminology sometimes differs from the terminology employed in ASOP No. 27. These definitions and the subsequent discussion presume a probability distribution of future investment returns. They therefore rely on an underlying reference portfolio (the portfolio in which plan assets are presumed to be invested over the measurement period, reflecting the intended asset allocation and rebalancing approach) to provide this basis for the probability distribution.

- *Average*: A statistic calculated from a sequence of values, which can be either historical returns or a single scenario of future returns. In other material, the word "average" is used to describe a calculation performed on a random variable. To avoid confusion, this practice note will use other terms to describe results that apply to random variables. The two types of average returns addressed by the practice note are:

    - *Arithmetic average return*: Calculated from a sequence of periodic returns by dividing the sum of the rates of return by the number of periods. For example, the arithmetic average of 2%, 5%, and -1% is (2% +5% -1%) ÷ 3 = 2%.

    - *Geometric average return*: Calculated from a sequence of periodic returns by first converting each of them to the amount that would be accumulated during the period from an investment of $1. For example, the single period accumulation that corresponds to a 10% return is 1.1, while the accumulation corresponding to a -5% return is 0.95. The geometric return over N periods is determined by taking the Nth root of the N periodic single period accumulations and subtracting 1 from the result. For example, the geometric average of 2%, 5%, and -1% is $(1.02 \times 1.05 \times 0.99)^{1/3} - 1 = 1.97\%$. As discussed in more detail later, if returns vary from one period to the next, the geometric

average return over multiple periods will always be less than the arithmetic average return.

- *Terminal wealth*: The amount that accumulates from an initial investment of $1. For any value of terminal wealth at the end of N periods, the equivalent discount rate is determined by taking the Nth root of terminal wealth and subtracting 1.

- *Independent and identically distributed (IID):* In probability theory and statistics, a sequence of random variables is independent and identically distributed if they share the same probability distribution and each variable is independent of all others. That is, each random variable is unaffected by the variables that came before. The assumption that observations will be IID tends to simplify the underlying mathematics of many statistical methods. The assumption is important in the classical form of the central limit theorem, which states that the probability distribution for IID variables with finite variance approaches a normal distribution. Not all actuaries consider the assumption of IID to be an adequate representation of projected investment returns.

The following two terms describe properties or results developed from the probability distribution of a random variable, such as the output from a stochastic simulation:

- *Mean or Expected value*: The average of possible values for a random variable weighted by the probability associated with each value. In stochastic analysis this outcome is estimated to be the average of the variable in question for all simulated scenarios.

  The word "expected" is often used in other contexts to refer to a single outcome that is considered likely. For example, an individual might say that the home team is "expected" to win a game in which it is favored although a loss is possible. Because its usage in this sense is common, this practice note instead generally refers to "mean."

  Some sources will describe average returns developed from historical results or a single sequence of forecast outcomes as mean returns. (For example, they may use phrases such as "arithmetic mean return" or "geometric mean return.") This practice note uses "mean" only to describe a statistic related to a random variable, not a statistic calculated from a sequence of values.

- *Median*: A value that separates the upper 50% from the lower 50% of the distribution of outcomes for a random variable.

The arithmetic and geometric average returns and the terminal wealth outcomes are themselves random variables. Statistics such as the following may be useful in determining a basis for setting an investment return assumption:

- the mean value of arithmetic average return (ASOP No. 27 refers to this as "forward-looking expected arithmetic return.")

- the mean value of geometric average return (ASOP No. 27 refers to this as "forward-looking expected geometric return.")

- the mean and median values of terminal wealth

- the equivalent discount rates associated with the mean and median values of terminal wealth

II. EXAMPLE

Much of the discussion that follows will consider these calculations as applied to a set of simulated future capital market outcomes such as those developed from a stochastic forecast. These outcomes can be presented in a table of results, arranged with each scenario as a row and results for each simulation year as a column. The analysis of historical results or a deterministic forecast would, in contrast, entail only one set of outcomes.

Exhibit 1

	Annual Return							
	Year					*Arithmetic Average*	*Geometric Average*	*Terminal*
Scenario	*1*	*2*	*3*	*4*	*5*	*Return*	*Return*	*Wealth*
A	5%	16%	20%	7%	-4%	8.8%	8.5%	1.50
B	14%	1%	6%	-12%	3%	2.4%	2.0%	1.11
C	1%	14%	26%	-3%	18%	11.2%	10.7%	1.66
D	22%	-4%	6%	11%	-3%	6.4%	6.0%	1.34
E	6%	14%	-3%	-8%	12%	4.2%	3.9%	1.21

The statistics for each scenario are determined as described above. For example, the arithmetic average for scenario A is equal to (5% + 16% + 20% + 7% - 4%) / 5 = 8.8%. The geometric average for the same scenario is $(1.05)(1.16)(1.20)(1.07)(0.96)^{1/5} - 1 =$ 8.5%. Similarly, the terminal wealth is (1.05)(1.16)(1.20)(1.07)(0.96) = 1.50.

The combination of model-generated scenarios makes up a collection of random variables for which additional statistics can be calculated. The mean and median of arithmetic average, geometric average, and terminal wealth are shown below. The equivalent discount rates that generate terminal wealth figures are also calculated.

Simulation results	*Mean*	*Median*
Arithmetic average	6.6%	6.4%
Geometric average	6.2%	6.0%
Terminal wealth	1.36	1.34
Discount rate associated with terminal wealth	6.4%	6.0%[2]

Reporting historical returns
Over a single investment period, arithmetic and geometric calculations of return are equal by definition. For multiple periods, however, the average returns will be equal only if each of the time-period returns are the same. To the extent that there is return volatility, the arithmetic average will be higher than the geometric average return, as the above example illustrates.

Standards have been developed specifically for use in performance reporting. These require linking investment performance over multiple periods geometrically, not arithmetically. This approach produces the single rate of return that would have produced the same rate of growth as the known but varying sequence of past returns. Suppose, for example, that the sequence of returns illustrated in scenario D actually came to pass. In that case, the terminal wealth would reconcile with the geometric average of the portfolio returns in that scenario: $(1.060)^5 = 1.34$.

The selection of a return assumption for discounting future cash flows is a different exercise. The convention that has been established for performance reporting may not necessarily be the most desirable when calculating liabilities. This practice note presents issues for the actuary to consider when determining which approach best fits the purpose.

III. FORECAST MODELS—THE EFFECT OF UNCERTAINTY

The analysis of past performance does not consider uncertain future outcomes, but forward-looking/forecast models typically do, and such analysis is critical to actuarial work. Intuitive conclusions based on analysis of historical results may not apply to the probability distributions of future returns.

Consider this highly simplified example: A distribution of outcomes based on only two potential return outcomes, +20% or -10%, with a 50% probability assigned to each. The returns for each year are presumed to be independent, without any serial correlation or reversion to mean.

[2] Note that the median geometric average return equals the discount rate equivalent of median terminal wealth by definition.

Exhibit 2

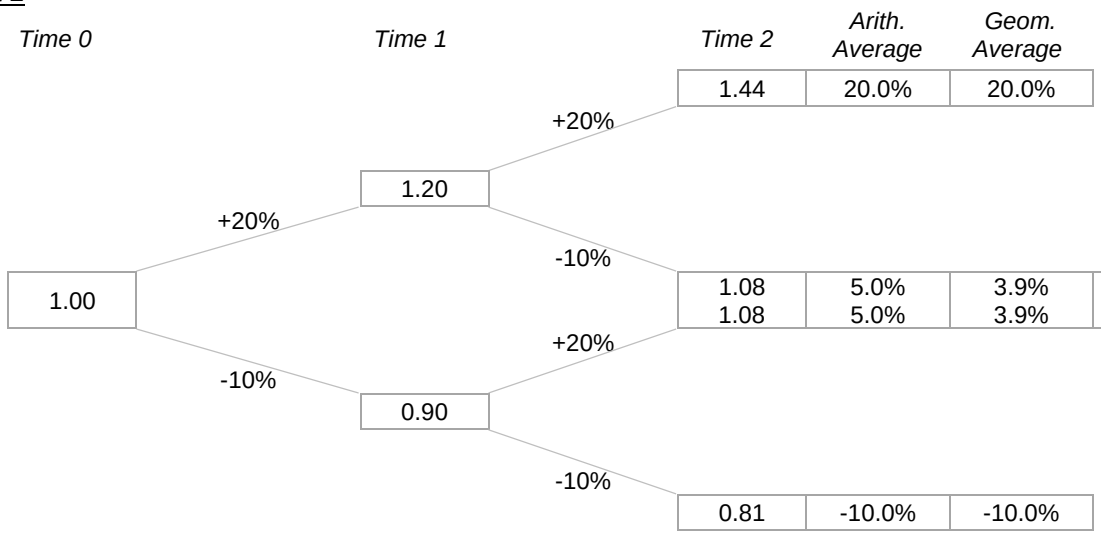

The mean of each year's return is, of course, 5%. The median annual return result is also 5%.[3] The annual returns are symmetric in the sense that the median equals the mean.

Even though the distribution of annual returns is symmetric, the distribution of terminal wealth in two years is not. The highest terminal wealth (1.44) exceeds the median (1.08) by more than the median exceeds the lowest outcome (0.81). The same observation applies to the geometric average return but not to the arithmetic average return. When considering the distributions of geometric average return or terminal wealth, the mean outcome will exceed the median outcome. (If returns are constant, the two statistics will be identical).

	Mean	*Median*
Arithmetic average	5.0%	5.0%
Geometric average	4.5%	3.9%
Terminal wealth	1.10	1.08
Discount rate associated with terminal wealth[4]	5.0%	3.9%

IV. RELATIONSHIPS AMONG STATISTICS

The relationships among these statistics are easiest to evaluate when future years' distributions of returns are considered to be IID, as was presumed in the example above. While this assumption forms the basis of many statistical models and conclusions, it may not incorporate the dynamics of observed return patterns.

[3] In this example, half of the returns are +20% and the other half are -10%. It is typical to calculate the median of an even number of outcomes as the midpoint of the middle two outcomes.

[4] The single rate that reproduces the mean or median value for terminal wealth. For example, the mean terminal wealth of 1.10 would be generated by a constant 5% annual return.

Nonetheless, the IID assumption allows for straightforward application of statistical concepts and permits the representation of portfolio return as a lognormally distributed random variable. (See Appendix 1 for additional discussion of the lognormal model.) This assumption facilitates the demonstration of certain numerical relationships that will be discussed further below. Note that the relationships may be valid even when prospective returns are not IID; at least some of these same relationships will be found in the output from any scenario generation model when applied over sufficiently long periods of time.[5]

Arithmetic average and geometric average returns:

- Over a single period, arithmetic and geometric measures of return are identical by definition.
- Over multiple periods of returns, either historical or projected to occur over a single trial, the arithmetic average return will equal the geometric average return only if all periodic returns are equal. If there is any return volatility, arithmetic average return will exceed geometric average return. Over multiple trials, the mean arithmetic average return will therefore exceed the mean geometric average return.
- The mean of the distribution of geometric average returns will tend to decrease as the projection period increases (given some level of return volatility). There are a number of estimates for the relationship between mean arithmetic (A) and mean/median geometric average (G) returns over long time horizons. The most common approximation, although not the most accurate, is $G \approx A - \text{Variance}/2$, where variance is that related to single-period returns.[6]

Arithmetic average return and terminal wealth:

- The mean of the distribution of arithmetic average returns relates to mean terminal wealth. In other words, accumulating assets at the mean arithmetic average rate is expected to produce the mean terminal wealth.

Geometric average return and terminal wealth:

- The median of the distribution of geometric average returns corresponds to median terminal wealth. Also, because mean geometric average return converges to median geometric return as the projection period increases, mean geometric average return also ultimately equates to median terminal wealth.

[5] The wide range of possible simulation techniques complicates efforts to draw definitive conclusions about the circumstances under which various relationships will be exhibited. Analysis of the outcomes that result under IID conditions is relatively straightforward, but a broader class of simulations will also exhibit these relationships. IID properties should be viewed as sufficient but not strictly necessary to produce these results. Empirical analysis of simulated results may be the most effective way to assess various statistical relationships.

[6] For example, in Exhibit 2, the 1-period standard deviation of returns is 15%, the arithmetic average A=5% and the geometric average G=3.9% and the approximation holds true: $3.9\% \approx 5.0\% - 15\%^2/2$. See the referenced Mindlin papers for a more complete discussion of this formula, along with an array of alternative estimation approaches.

An actuary referencing forecast results might need to review and test the distribution of outcomes from a particular capital market model to determine how well various relationships hold. In other words, the actuary might need to evaluate, rather than presume, connections such as the critical linkage between the mean values for arithmetic average and terminal wealth.

V. ANALYSIS OF FORECAST RETURNS

The actuary's determination of an expected return assumption might be based on simulated future capital market outcomes along with, or in place of, a review of actual/historical capital market results. A stochastic forecast model will generate an array of possible results that can be characterized as arithmetic or geometric average returns. Some characteristics associated with each statistic may be of interest, including:

No expected gain/loss
This is a traditional actuarial objective. If the expected return assumption is set equal to the discount rate equivalent of mean terminal wealth, the expected gain or loss on assets in the future, in dollar terms, will be zero. Appendix 3 of ASOP No. 27 asserts that the mean arithmetic average return (forward-looking expected arithmetic return) will produce no expected gains or losses. This result would be anticipated from a model based on IID-type parameters, but may not be found in other models that incorporate implied mean reversion.[7] In such cases, it might be appropriate to determine the discount rate equivalent of the mean terminal wealth result rather than to approximate that outcome by use of the arithmetic average.

As indicated above, the mean geometric average return converges to the discount rate that corresponds to median terminal wealth. In other words, gains and losses will occur with equal frequency when measured with respect to the mean geometric average return. The magnitude of the gains, however, will typically not be the same as the magnitude of the losses. The gains associated with high outlier outcomes will generally exceed the losses associated with low outlier outcomes.

If the assumed expected return is set to the expected geometric average and that geometric average return is realized over a given period, no gain/loss will result. If the assumed return is set to the expected arithmetic average return and that arithmetic average return is realized, however, there is likely to be a loss. Unless the return is realized as a constant rate every year, a loss will arise. The geometric average return that is realized, which corresponds to the accumulation of wealth, will be less than the arithmetic return (see Section II). Because the experienced return amounts will almost certainly not be returned as a constant rate, an arithmetic average return that is *greater* than the investment return assumption must be realized in order to avoid a loss.

[7] In models with mean reversion tendencies, the mean arithmetic average return result is likely to exceed the discount rate equivalent of mean terminal wealth. This imbalance arises from such models' tendency to pull outlier results within a given sequence of simulated returns back toward the median over the successive years. Doing so effectively suppresses "longitudinal" volatility (the range of accumulated wealth outcomes over time) while leaving "cross-sectional" volatility (the range of return outcomes for any one simulation year) unaffected.

For example, consider annual returns based on a distribution with a mean of 6%. Returns consistent with this distribution will, at the median, result in an arithmetic average return of 6% but a geometric average return below 6%. If the actuary uses a constant rate of return of 6% as the valuation assumption, then the median outcome will produce losses relative to this assumption. It will take an above-median set of returns to produce a result that is consistent with a constant 6% annual return. Historical performance is thus conventionally measured with geometric averages, not arithmetic.

Credibility/robustness[8]
The mean of a random variable is much more sensitive to outlier values than its median, because the mean value is affected by the existence of a few large outlier values, while the median is not. Because geometric average return corresponds to the median terminal wealth statistic, it is considered to be a more robust outcome from a capital market simulation model than is the arithmetic average return.

This characteristic becomes especially important if the actuary believes that outlying scenarios in a probability distribution are not fully credible. Certain statistical techniques may also be considered to address this situation. For example, the outlying scenarios may be truncated, or their values may be replaced with threshold values. It may be necessary to consider the specific situation, including the purpose of the measurement, before making any such adjustments to the distribution.

Conservatism
Because mean arithmetic average return will almost always exceed mean geometric average return (and will never be less than it), the use of the arithmetic average for discounting purposes would be viewed as a less conservative assumption.

VI. CONSIDERATIONS FOR ACTUARIES

As noted earlier, the geometric average of historical returns is the single rate that would have generated the same wealth accumulation as actually observed. Reference to historical results when setting assumptions about the future raises additional considerations. The likelihood of similar outcomes recurring is affected by differences between current economic conditions and those observed in the historical period analyzed. Simply using historical return averages as estimates of future returns will generally not capture the effect on future returns of key drivers such as current inflation levels, interest rates, and stock market valuations.

---

[8] These terms are related in the sense that they connect to the level of confidence that might be attributed to a given modeling result.

-- The term *robustness* relates to (1) the sensitivity of a given result to outlier data in the distribution from which it is derived, and (2) the ability of a test or result to provide valid insight even if the model presumptions are altered or violated.

-- The term *credibility* as employed in this context relates to the level of reasonableness/validity associated with a given simulation result; it seems rational to assert that reliance on a less-robust forecast result would be considered less predictive of actual future outcomes.

Note that even if a forecast model were calibrated to fully align with historical results—asset class means, standard deviations, and correlations that exactly match historical statistics—it would still produce a range of outcomes rather than the deterministic/single outcome represented in the historical record.

The generation and calibration of economic scenarios involves a host of decisions, and at least some simplification is generally necessary. The effect of these simplifications is an important consideration when assessing the credibility of simulated results. For example, cyclical qualities of capital markets might not be accurately simulated in modeling. One view is that mean-reversion tendencies exist in capital market outcomes over time.[9] A model that does not incorporate a mean reversion quality—e.g., one based on an IID presumption for the generation of annual outcomes—would be expected to produce a range of outcomes that is broader than a model that does reflect mean reversion. Because mean wealth outcomes are disproportionately affected by high outlier results, the actuary might consider the plausibility of return/wealth outcomes that are heavily dependent on such high outlier results.

One might also consider whether to focus on a mean outcome or on a distribution of outcomes as the basis for decision-making. When considering events that are repeatable, gains from one iteration are available to offset losses that occur in other iterations. For example, consider a bet of one dollar on the selection of a single integer from 1 to 1,000 with a payoff of 1,000:1. The expected value of this wager is one dollar. In this case, the highly likely but relatively small losses might be expected to offset the relatively unlikely but very large gain associated with a win. As long as the one-dollar bet is a small portion of the bettor's overall wealth, the game can be repeated often enough that the few favorable outcomes can be expected to offset the effect of the more numerous unfavorable outcomes.

However, if the number of expected incidences of betting is reduced for any reason (e.g., the bet is a large portion of the bettor's wealth), the situation changes. If there will be only a few betting opportunities, it might be more appropriate to focus on the distribution of expected outcomes, with greater focus on likely as opposed to mean outcomes. This recognizes that gains from the improbable but extremely favorable outcome are unlikely to be available to offset losses from the much more probable unfavorable outcomes. Of course, this does not necessarily imply that one should focus on the midpoint of the distribution of outcomes. Depending on objectives, a 50% chance of achieving the targeted result may or may not be sufficient.

[9] Mean-reversion tendencies would presumably result from constraints on the range of economic activity and capital market results, e.g., those imposed by resource/workforce/productive capacity limitations in the overall economy, current or simulated levels of interest rates vs. presumed normative levels, the level of equity pricing in comparison to historical mean price levels, and through the operation and underlying objectives of government fiscal and monetary policies. Note that the efficient market hypothesis implies that prices follow a random walk and consequently that rates of return are IID.

Similarly, although a distribution of outcomes may be developed (mathematically or through a simulation of outcomes),[10] there will ultimately be only one outcome. Gains from other favorable simulations will not be available to offset losses from unfavorable realized results. Thus, averaging the results from an array of potential outcomes may result in a measure that has limited practical value, especially in situations where it is more likely that actual experience will fall short of that average outcome. For this reason, a focus for decision-making might be on the distribution of results, such as the median and various percentile outcomes.

VII. CONCLUSIONS

The conclusions from this practice note can be briefly summarized as follows:

- When evaluating historical return statistics, the use of geometric average return results is generally appropriate.

- When analyzing simulated future outcomes to select an expected return assumption to use as a discount rate, consideration may be given to both mean geometric and arithmetic average results, along with other related statistics such as the discount rate equivalent of mean or median terminal wealth.

- The actuary might expect that the use of an assumption based on the mean arithmetic average, or the return rate that generates the mean terminal wealth outcome, will produce no expected future gain or loss.[11] However, the gain/loss parity results from the greater dollar gain associated with high outlier outcomes vs. the smaller loss associated with low outlier outcomes. Thus, despite there being no gain or loss *on average*, the use of this assumption actually involves a greater-than-50% chance of a loss being incurred.

- In the context of simulated future results, over long periods the mean geometric average will align with the median wealth outcome, thus balancing the expected likelihood of gains and losses in the future. The mean geometric average is less sensitive to the influence of outlier results than is the arithmetic average, which means that it is the more robust outcome from capital market modeling.

- These conclusions are relevant primarily to the use of the investment return assumption as a discount rate in the measurement of liabilities. Appendix 3 describes two situations in which other considerations may apply.

[10] The discussion in this paper focuses on the distribution of uncertain future outcomes. An alternative framework would take the financial commitment of a pension fund as a given and then derive a distribution of present values consistent with this commitment. This concept is discussed in the Mindlin paper. Either approach may present a useful framework for decision-making.

[11] As noted earlier, the presumed equality in these two forecast outcomes might not be found in models that incorporate significant mean-reversion tendencies; i.e., calculated mean arithmetic average returns might exceed the level implied by mean terminal wealth.

The implications of using investment return assumptions based on arithmetic or geometric returns are surprisingly complex. An actuary considering the selection of an investment return assumption for discounting over long periods of time may find it helpful to consider the issues and concerns raised in this practice note.

APPENDIX 1

Derivation of Conclusions From IID Assumptions

The relationships among statistics are easiest to evaluate when future years' distributions of returns are considered to be independent and identically distributed (IID). While this assumption forms the basis of many statistical models and conclusions, this treatment is a simplification in that it does not incorporate some dynamics of return patterns actually observed.

Statistical models based on IID principles, however, have some theoretical basis and exhibit a number of useful and noteworthy relationships. The weak form of the Efficient Market Hypothesis implies that stock prices do not depend on the past prices and will instantly react to new information. This implies that successive returns (annual or instantaneous) are independent random variables. If we also assume that the instantaneous, continuously compounded rates of return are independent, identically distributed (IID) random variables, then stock prices will have a lognormal distribution. This model forms the basis of the capital asset pricing model, the Black-Scholes model, and other widely referenced models. The discussion of modeling stock prices with lognormal distributions in this appendix is based in part on Chapter 18 of *Derivatives Markets* (Third Edition), 2013, by McDonald, R.L., Pearson Education.

A key reason we focus on the continuously compound rate of return as opposed to the annually compounded return is that in order to use the Central Limit Theorem (CLT), we must take the average of a sequence of random variables and annual returns are compounded, not averaged. Converting to continuously compounded return allows us to take an average in the exponent and thus use the CLT. In particular, if we divide the interval from *[0,1]* into *n* equal time period of length *1/n*, and assume random annual rates of return $r_1, r_2, \ldots, r_n$, then a stock with price S_0 at time *0* will have price at time *1* of

$$S_1 = S_0*(1+r_1)^{(1/n)}(1+r_2)^{(1/n)}\ldots(1+r_n)^{(1/n)},$$

which does not simplify easily. However, if we convert each r_i to a continuously compounded $\delta_i = \ln(1+r_i)$, the stock price at time *1* will be

$$S_1=S_0 * \exp(\delta_1/n)*\exp(\delta_2/n)*\ldots*\exp(\delta_n/n)= S_0 *\exp(\textstyle\sum_{i=1}^{n} \delta_i/n).$$

Note that the term in the final exponent, $\sum_{i=1}^{n} \delta_i/n$, is the average of the continuously compounded rates of return. We then apply the Central Limit Theorem to the exponent and see that $\sum_{i=1}^{n} \delta_i/n$ converges to a normal distribution as *n* goes to infinity, provided that the mean and variance of each δ_i are the same finite constants.

If $\sum_{i=1}^{n} \delta_i/n$ converges to $N(\mu,\sigma)$, then $S_1 = S_0 * \exp(N(\mu,\sigma))$ has a lognormal distribution. In this case, the random 1-period annual rate of return $r = S_1/S_0 - 1$ will have the following properties:

$$\text{Mean: } m = exp(\mu+\sigma^2/2) - 1$$
$$\text{Median: } e^{\mu} - 1$$
$$\text{Variance: } s^2 = exp(2\mu+2\sigma^2) - (exp(\mu+\sigma^2/2))^2$$

Note that the median is below the mean and the difference is approximately half the variance as noted previously. Alternatively, given the arithmetic annual return *m* with standard deviation *s*, we can solve for continuously compounded lognormal parameters μ and σ as follows:

$$\sigma = sqrt(ln(s/(1+m))^2+1)$$
$$\mu = ln(1+m) - \sigma^2/2$$

For a projection covering N investment periods, mean arithmetic average return, mean geometric average return, and the discount rate equivalents of mean and median terminal wealth may be calculated directly. If we assume the expected annual rate of return is $m = 5\%$ and the standard deviation is $s = 15\%$, under the lognormal model, the continuously compounded parameters are $\mu = 3.87\%$ and $\sigma = 14.21\%$. The resulting median is $exp(3.87\%)-1 = 3.94\%$. Those statistics are shown in the graph in Figure 1, and exhibit the following relationships:

- Mean arithmetic average return is constant (independent of *N*) and is equal to the expected or mean value of the single period return.

- Mean geometric average return equates to the arithmetic average for a single-year period, and then decreases over time (as *N* increases) to the median.

Figure 1[12]

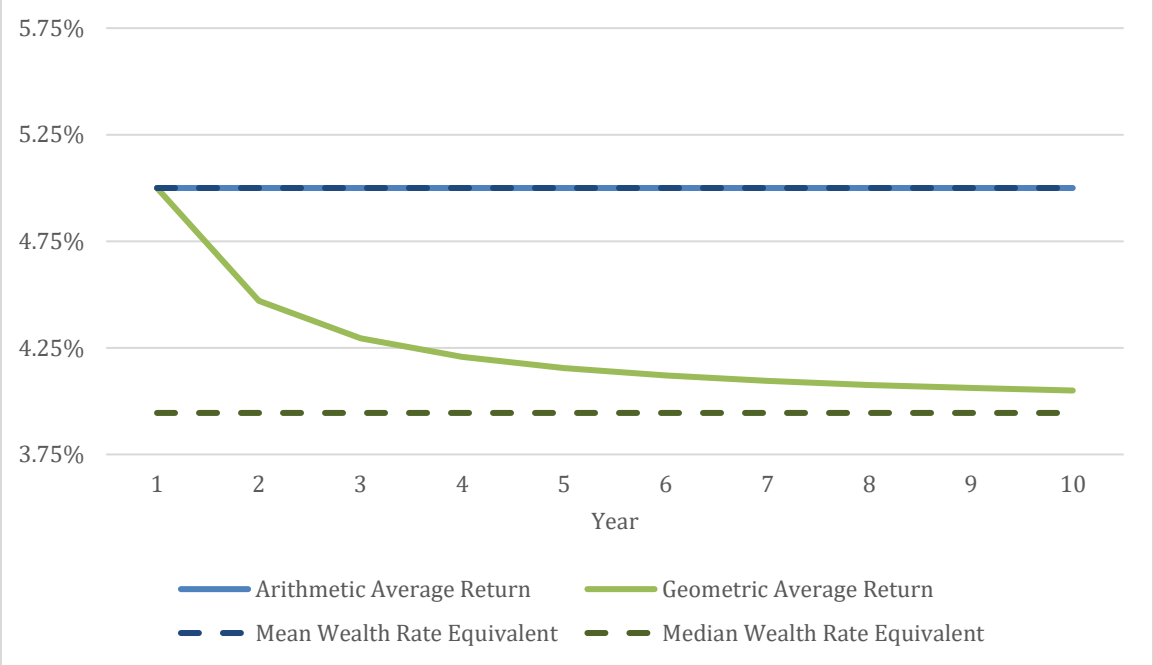

Terminal wealth

The objective in pension plan funding is not to achieve a particular level of investment return, but rather to accumulate an amount over time that is sufficient to provide for the payment of pension obligations. For that purpose, the most relevant statistics are those that relate to wealth accumulation, and similarly, the equivalent discount rates corresponding to those wealth statistics. In the simplified statistical model, these statistics will exhibit the following characteristics:

- Mean terminal wealth has an equivalent discount rate that is constant independent of N, and equates to mean arithmetic average return.

- Median terminal wealth has an equivalent discount rate that, by definition, equates to median geometric average return.

Mean geometric average return decreases over time as N increases; over long projection periods, it asymptotically approaches the equivalent discount rate that equates to median terminal wealth.

Relationships Referenced in ASOP No. 27 – Appendix 3

Some expected relationships between various statistical outcomes are referenced in ASOP No. 27, Appendix 3. These references are essentially the same as those quoted above, i.e., statistical connections that an actuary would expect to see in statistically based models incorporating IID-type principles.

[12] Results of a return simulation based on IID presumption, lognormally distributed returns, 5% mean return, and 15% standard deviation.

In particular, the Appendix references two key expected relationships and, as noted, uses somewhat different terminology than is employed in this practice note:

- The use of a forward-looking expected geometric return as a discount rate will produce a present value that generally converges to the median present value as the time horizon lengthens (i.e., if the actuary determines a funding obligation using the forward-looking expected geometric return to discount the obligation to produce a present value, it is expected that in the limiting case there will be enough money to fund the obligation 50% of the time).
- The use of a forward-looking expected arithmetic return as a discount rate will generally produce a mean present value (i.e., there will be no expected actuarial gains and/or losses).

APPENDIX 2

Implications of Assumptions Other Than IID

Actuaries may also use more complex capital market/forecast models that do not adhere to the IID convention. Many models have provisions to address differences between initial capital market conditions and "normative" conditions; e.g., current interest rates may be considered lower than the long-term norm and thus future rates will have a tendency to rise. Similarly, equity valuations could be viewed as out of sync with long-term valuation levels and have a tendency to rise or fall over time to compensate.

In addition to trends related to initial-normative capital market conditions, some models may also incorporate tendencies toward mean reversion within the generated scenarios, which implies that when return results in a given scenario are simulated to fall extremely far from the normative trend, those extreme outcomes will have a tendency to be reversed over time. For example, extremely favorable equity returns may be presumed to imply levels of economic growth, P/E ratios, and utilization of workforce, resource, and production capacities that are higher than normal. Given modeled constraints on these parameters, the result may be a bias toward unfavorable equity returns in successive periods that act to suppress prospective returns and push accumulated results closer toward the more typical range. Similarly, simulated high fixed-income returns generally result from decreases in yields that will tend to be reversed over time.

These types of model characteristics will tend to disrupt some of the relationships that were evidenced in the simpler statistical model reviewed in Appendix 1, as illustrated in Figure 2.

Figure 2[13]

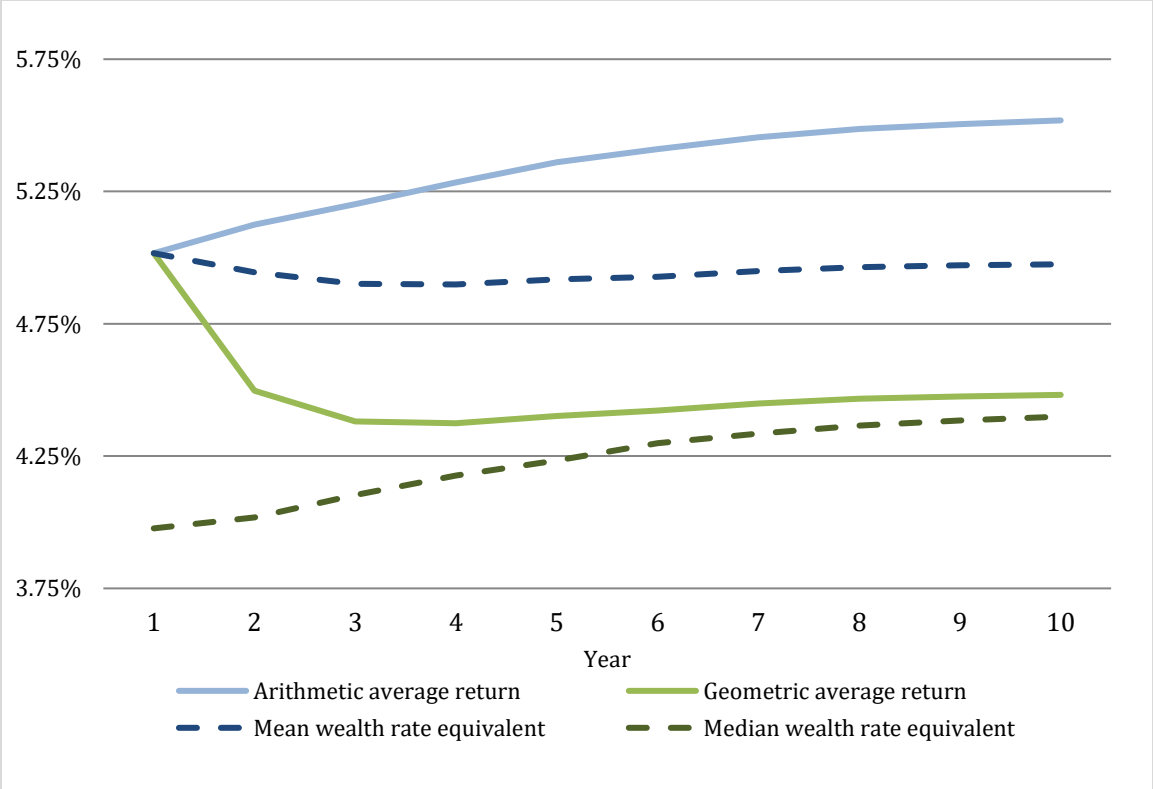

As the above example illustrates, results from more complex models may create disconnects in at least two critical relationships:

- a trend in rates rather than constant emerging rates for mean arithmetic average and mean terminal wealth; and
- a gap rather than equality between emerging results for mean arithmetic average and mean terminal wealth.

The first outcome is a result of the tendency for initial capital market conditions to revert to normative levels over time. The second outcome is caused by the tendency for mean reversion within the capital market simulation, so that the emergence of extremely high or extremely low return/wealth outcomes creates a tendency for offsetting outcomes in successive periods—which acts to pull extreme wealth outcomes back toward median levels.

[13] Results of a return simulation based on normally distributed returns, 5% mean return in year 1 grading up to 6% mean return in year 7, and 15% standard deviation. The model also incorporates some reversion of simulation outcomes to the mean, so that a high percentile outcome in a given year leads to an increased likelihood for a low percentile outcome in a successive year; this acts to suppress the distribution of accumulated returns (e.g., the gap between mean and median wealth outcomes). In this particular model, 50% of the difference in any single year between the simulated scenario return and the mean return is applied as an adjustment to future returns over the subsequent five years in a declining pattern. While this model may be simpler than those typically used in practice, it is sufficient to illustrate the effects of mean reversion on the forecast outcomes shown above.

APPENDIX 3

Special Considerations for Specific Applications

The body of this practice note focused on idealized situations in which investment return assumptions are used as discount rates and applied over multiyear periods. Other applications of investment return assumptions introduce elements not yet considered. In these applications the concepts underlying this analysis generally remain relevant, but the specific conclusions may need to be adapted. This appendix considers two such circumstances.

- *Application of Return Assumption in U.S. Corporate and Plan Accounting*
  U.S. corporate accounting rules applicable to sponsors of pension and retiree welfare plans call for a different application for an investment return assumption. The expected return on asset assumption represents a long-term expectation, described in Accounting Standards Codification (ASC) 715 as "the average rate of earnings expected on the funds invested or to be invested to provide for the benefits included in the projected benefit obligation." But it is applied to a current value of assets to calculate the expected return on asset amount that is a component of benefit cost for the current year.

  Because the rate is applied to a current asset value to develop a single-year return amount, its application is quite different than for an investment return assumption that is applied as a constant discount rate or used as an estimate of asset accumulation over longer time horizons. It is only when a return assumption is applied over multiple time periods that many of the issues discussed earlier arise. A mean arithmetic average return figure is arguably more compatible with the function of estimating a single year's investment return.

  Plan sponsors may be required to prepare financial reports according to several accounting standards. The considerations of this practice note may apply differently to each. For example, while the provisions of ASC 715 referenced above (covering the determination of pension expense) call for an investment return assumption to generate an estimate of single year return, the investment return assumption defined under ASC 960 (covering plan accounting) is employed as a discount rate. One might therefore reasonably set the ASC 715 rate based on the arithmetic average return and the ASC 960 rate based on a geometric average return for the same plan.

- *Gain-Sharing Features*
  In some plans, the level of benefits provided may depend on asset performance. If the actuary is charged with determining a rate of return that is allocable to the payment of already-defined benefits under the plan, these features may have an impact on the return assumption that is considered appropriate for that purpose.

  For example, some public sector and Taft-Hartley plans may have gain-sharing

arrangements where ad-hoc COLAs or 13th checks are granted based on return performance above a certain threshold. These features create an asymmetry in the distribution of net returns available to pay already-defined (primary) benefits.

Similarly, provisions such as minimum or maximum interest crediting to individual accounts or other embedded options that convert a portion of investment returns to a benefit may also alter the distribution of returns available for funding the primary benefits.

Valuation of such features is beyond the scope of this practice note. Additional analysis of the net return available for funding primary plan benefits may be warranted.

These circumstances are only representative; others may also be encountered when applying expected return assumptions in practice. These situations may not lead an actuary to a definitive conclusion that either an arithmetic or geometric average return is appropriate. Some considerations may support one statistic while others support another. In such circumstances, an actuary may consider a modified rate representing a blend of the two statistics to be a reasonable choice.

SUGGESTED REFERENCES

Derivatives Markets (Third Edition)
McDonald, R.L.
Pearson Education
2013

"Geometric or Arithmetic Mean: A Reconsideration"
Jacquier, Eric, Alex Kane, and Alan J. Marcus
Financial Analysts Journal
November/December 2003

Global Investment Performance Standards
CFA Institute
2010

Investment Science
Luenberger, David G.
Oxford University Press
1998

Geometric Return and Portfolio Analysis
McCulloch, Brian
New Zealand Treasury Working Paper 03/28
December 2003

Present Values, Investment Returns and Discount Rates
Mindlin, Dmitry
CDI Advisors LLC
May 2013